

Closing the loop in model-based wind farm control

Doekemeijer, B.M.

DOI

[10.4233/uuid:35022242-b3ba-48d9-b3f6-c6ab20b8cc19](https://doi.org/10.4233/uuid:35022242-b3ba-48d9-b3f6-c6ab20b8cc19)

Publication date

2020

Document Version

Final published version

Citation (APA)

Doekemeijer, B. M. (2020). *Closing the loop in model-based wind farm control*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:35022242-b3ba-48d9-b3f6-c6ab20b8cc19>

Important note

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

Copyright

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

Takedown policy

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

CLOSING THE LOOP IN MODEL-BASED WIND FARM CONTROL



PHD DISSERTATION OF
BART M. DOEKEMEIJER

Closing the loop in model-based wind farm control

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op 16 september 2020 om 15:00 uur

door

Bart Matthijs DOEKEMEIJER

Ingenieur op het gebied van systeem- en regeltechniek,
Technische Universiteit Delft, Nederland,
geboren te Dordrecht, Nederland.

Dit proefschrift is goedgekeurd door de promotoren:

Prof. dr. ir. J.W. van Wingerden
Prof. dr. ir. M.H.G. Verhaegen

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. ir. J.W. van Wingerden	Technische Universiteit Delft, promotor
Prof. dr. ir. M.H.G. Verhaegen	Technische Universiteit Delft, promotor

Onafhankelijke leden:

Prof. dr. ir. C.J. Simão Ferreira	Technische Universiteit Delft
Prof. dr. ir. D.A. von Terzi	Technische Universiteit Delft
Prof. dr. ir. L.Y. Pao	University of Colorado Boulder, VS
Dr. ir. P.A. Fleming	National Renewable Energy Laboratory, VS

Overige leden:

Dr. ir. P.M.O. Gebraad	Siemens Gamesa Renewable Energy, Denemarken
------------------------	---



Keywords: closed-loop wind farm control, wake steering, yaw-based wake redirection, power maximization, model-based optimization, real-time model adaptation, nonlinear state estimation, floris, wfsim, sowfa

Printed by: Gildeprint B.V., Enschede, Nederland

Copyright © 2020 by B.M. Doekemeijer

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

ACKNOWLEDGEMENTS

Pursuing a doctorate at the Delft University of Technology is often said to be a long, exhausting and demotivating journey. On more than one occasion, I have spoken with PhD students that were on the verge of quitting due to poor or inconsistent supervision by their promotor, an uninspiring topic and not in the least the immense pressure to publish mixed with the uncertainty of what research topics to pursue and whether those will lead to fruitful results. Fortunately, my experience has been nothing like these horror stories and these four years have actually been the happiest four years of my life. This PhD program has allowed me to learn, explore, travel, connect and grow as a person, both on a professional and a personal level. This is mostly by virtue of a number of people, whom I wish to thank here.

Firstly, I want to thank Jan-Willem from the bottom of my heart for being the best supervisor I could ask for. Your door was always open in case I ran into issues, whether these were related to work or to my personal life. It is clear to see that you are passionate about your work and that reflects on the people around you. I admire your relentless optimism, your strong work ethic and your kind personality. Thank you for giving me the opportunity to do a PhD in your group and thank you for all the freedom and trust, even allowing me to take the lead in our industrial project and in the CL-Windcon project. You have enabled me to grow more than I could possibly ask for.

Secondly, I would like to thank my (ex-)companions Sjoerd, Sebastiaan, Joeri, Daan, Yichao, Maarten and Atindriyo. Sjoerd, I appreciate the long conversations we had about wind farm control, research at large, the scientific world and about our personal lives. You have a hilarious sense of humor and I am blessed to have had you around for the initial three years of my PhD. I wish you all the best in France and hope to meet you again sooner than later. Sebastiaan, I loved our short interactions in the office, your positive attitude, and our collaborative efforts in solidifying and disseminating our novel cutting-edge contributions. I had a blast spending three weeks with you in the US and I hope we will stay in touch, wherever we end up working. Joeri, thank you for accompanying me through all the CL-Windcon project meetings and being a very productive and reliable colleague overall. Daan, thank you for all the laughs during our Master's and the PhD, your openness and your relaxed presence. It's always great to have you around. Yichao, you are a great colleague and you are a positive influence everywhere. We need more people like you. Maarten and Atin, though we have not spent a lot of time together, I am sure you'll be successful in your journey in the group and I wish you the best of luck. Delphine, Livia, Unai, Zhenyu, Jean, Zhixin, Folkert, Tomas, Sander, Coen, Amin, Cees, Max, Clara, Tian, Gabriel, thank you for the few interactions that we have had and I wish you all the best. I am glad to have had you all as my colleagues.

Thirdly, I owe much to the DCSC staff. Not in the least, I thank Erica and the DCSC

secretariat for bringing life to the research department; Marieke, Heleen, Kiran and Francy. Kim and Manon, I wish you both the best in your careers. I was fortunate to witness several visitors of our department, such as Leif, Steffen, Sara and Andreas. Each of you have sparked unique ideas in my mind and have contributed to this thesis, even if you may not have realized it. I would also like to extend my thanks to Roald, Nivas, Raymond, Ben, Hector and Marcus for their hard work and insightful discussions during their graduation projects under my supervision.

Fourthly, I want to thank all the members of the PhD council; Mascha, Eline, Ankur, Vilborg, Alina, Abhilash, Hongpeng, Zongchen, Costanza, Andrea and Pieter. I am proud to have worked together with you all in improving the life of PhD students at our faculty. Furthermore, I am blessed to have met Pradeep, Pavlo and the rest of the meditation group – you have inspired me to meditate, helping me maintain a peaceful mind.

Fifthly, I want to express my deepest gratitude to Lucy and her (ex-) lab members: Dan, Jake, Fiona, Misha, Nikhar, Arnold and Cassandra. Lucy, you truly inspired me to pursue research in the field of wind energy. You have been incredibly helpful, patient, kind and supportive throughout my Master's and my PhD. I could not have asked for a better supervisor (and host!). I cherish the memory of our ski trip and the dinners with the lab and your family. Dan, thank you for the hike to Green Mt. in February 2017 – this is probably still the most beautiful hike I have ever done to this date. I hope you can find an opportunity to visit Europe some time soon again, if politics do not stump it.

Sixthly, I want to thank the people at the National Renewable Energy Laboratory in Boulder, CO. Specifically, whenever I visited, Paul, Jen, Andy, Chris, Eric, Ryan, Katie, Nick, Tony, Alan, Eliot and Matt made me feel very welcome from the first day onward. Paul, I look up to you as a researcher and very much enjoyed our conversations in Boulder and in Delft. Jen, I appreciate your positive attitude, your modesty and your strong work ethic. And of course, thanks to all the ultimate frisbee players at NREL for the fun games on the Friday afternoons. Besides my direct colleagues from TUDelft, CU and NREL, I must thank Yusik, David, Torben, Stoyan, Ervin, Renzo, Irene, Filippo, Johannes, Pieter, Sachin, Bastian, Paolo, Stefan, Giancarlo, Mikel, Sugoï and Tom for the pleasant collaborations and discussions over the years.

Seventhly, I want to thank all the people that I met during my time in Boulder that made my experience living abroad unforgettable. Jessy, you became one of my best friends in a matter of months, getting caught up in all kinds of crazy adventures in the US. You have such a great and positive attitude, a kind personality, fun humor and I am blessed to call you my friend. Imran, you are an amazing friend and I look forward to seeing you whenever I come visit. Furthermore, I must thank my old roommates, Vitali, Dmitry, Madison, Raniero, Charlie and Aidana, for providing the best home a Dutchman lost in the US can wish for. I want to thank Michael, Heiko, Abdullah, Cody, Gautham, Andrea, Fares, Amar, Anton, Kay, Jack, Katya and Roman for all our shared adventures, dinners, talks and hikes. You guys have had a huge share in making me the person who I am today and I cannot thank you enough for that. Some of my happiest memories are from my time with you all. I also want to express my appreciation for the friends I have met across Europe playing a much outdated videogame: Stefan, Johan, Pål, Pol, Hakeem,

Leiv, David, Oliver, Max, Martí, Pierre, Guillaume, Jamie, and Ben.

Eightly, I must thank Aashish, Bhagyashri, Marcel and Rudra for all the late nights shared, the spontaneous (and foolish) decisions we make and for always being there when I need you the most. Thank you Brian, for being the kindest person I know and an invaluable friend for life. You have trusted in me more than I trusted in myself at times. Of course, I must acknowledge the *droeloeboys*, Ilias and Maurice, as being two of my most cherished friends. With you I have shared many of my most memorable moments, be it a blessing, be it a struggle. Most of them were the latter. Thank you Elmira for all the support you have given me over the years. Also, I want to express my gratitude to Frank, Bernard, Artur, Stijn, Arthur and Liam for motivating and pushing me to grow, mentally but also very much physically. It was a blessing to see and share the gym with you guys every morning. Frank and Bernard, I especially appreciate our deep conversations about life. Additionally, thank you Veli, Gijs, Martijm, Tim and Jonathan for being good friends throughout the years, even though we do not see each other quite often enough. Celine, Niels, Arianna and Mike, I look forward to sharing adventures with you in the future.

Ninthly, it should go without saying that my parents and siblings, Niels, Tim and Roos, have played invaluable roles in my journey. You were always there in times of need. Roos, I hope to have set a good example as a PhD candidate. I know you'll do great.

Finally, I am blessed to have met Bernice during my PhD. While we were in the same classroom from ages twelve to eighteen, I was too shy to say hi back then. I am so glad that we reconnected and that I did gather up the courage seven years later. Bernice, you make me strive to be the best version of myself. Thank you for all your love and support. I am excited to share a future with you.

SUMMARY

Climate change is one of the most imminent threats to society today. In 2016, 195 countries signed the Paris agreement, in which these countries commit to limiting the increase of the global average temperature to 1.5°C above pre-industrial levels. The transition from fossil fuels towards renewable energy sources plays an invaluable role in achieving this objective. Among the renewable energy sources, wind energy is considered one of the most potent contenders. This energy source is forecasted to become the European Union's single largest source of electricity by 2050. However, for financial and societal reasons, wind energy often remains infeasible compared to non-renewable energy sources. One strategy to improve its feasibility is through wind farm control, in which the turbines inside a wind farm are controlled in a coordinated fashion to achieve a common objective. Wind farm control has the potential to improve the overall performance of wind farms by increasing the annual energy yield, reducing the structural loads experienced by the turbines and enabling wind farms to provide ancillary services to the electricity grid.

Wind farm control solutions can leverage different control variables in pursuit of various objectives. This thesis focuses on yaw-based wake steering for power maximization. Yaw-based wake steering is a wind farm control concept in which the upstream turbines are purposefully misaligned with the inflow to displace the waked flow away from downstream machines, thereby increasing the net power production. Over the past decade, yaw-based wake steering has matured from an academic concept towards an industrial product. The approach was initially tested in high-fidelity simulation, after which it was tested in wind tunnel experiments and finally in the field. However, while two out of the three field experiments in the literature show situational increases in the power production, the concept is only tested for two-turbine interactions under specific inflow conditions. The question of how much wake steering increases the annual power production of a commercial wind farm thereby remains unanswered. The third field experiment in the literature entails wake steering for a full year of wind farm operation at a commercial site. While the authors show similar situational gains as what is found in the literature, the gain in the annual energy prediction is estimated to be negligible, thereby calling the success of the wind farm control solution into question. This thesis addresses this scientific gap. Accordingly, the overarching objective of this thesis is formalized as

“Maturing wind farm control technologies for power maximization in a model-based closed-loop framework towards real-world practical applicability”

The posed thesis objective is divided in three distinct research questions. The first research question is formulated as *“What are the limitations of the current open-loop wind farm control solutions?”* Accordingly, this thesis evaluates an open-loop wind farm

SUMMARY

control algorithm leveraging a steady-state surrogate wind farm model on a commercial, onshore wind farm in Italy. In the field experiment, two- and three-turbine interaction is looked into, where the upstream two turbines are misaligned with the inflow. The results show that wake steering has potential to increase the power production significantly, by up to 35% for two-turbine interactions and up to 16% for three-turbine interactions. However, erroneous wake steering is prevalent in the experiments and occasionally leads to performance losses. The current open-loop wind farm controllers do not yet consistently improve the energy yield in wind farms. Erroneous wake steering is hypothesized to mainly originate from surrogate model inaccuracies. One significant modeling error is the power-yaw dependency of the upstream turbines. Additionally, the farm is surrounded by complex terrain, which is not captured by the surrogate model. Moreover, some of the wind turbines in the wind farm have a different hub height, of which the effects are not validated in the surrogate model. Finally, the inherent time-varying nature of the flow and the wind turbines is not accounted for in the employed steady-state approach.

In view of these results, this thesis defines two additional research questions. The second research question is formulated as “*Can we establish resilience to surrogate model uncertainty in modern wind farm control solutions through real-time measurement feedback and if so, how?*” Considering the losses seen in the field experiment, this thesis focuses on transitioning to the closed-loop architecture – using measurements to estimate the inflow wind conditions and adapt the surrogate wind farm model in real time to accurately capture the current wind farm behavior.

A novel theoretical framework is defined in this thesis that provides a quantitative measure for the *observability* of the farm – indicating how well the ambient conditions can be reconstructed from the available measurements. The analysis reveals which measurements are necessary to accurately reconstruct the ambient conditions using steady-state wind farm models. The results confirm the notion that local wind speed and wind direction estimates at each turbine are necessary for the estimation of a wind-farm-wide wind speed and wind direction, respectively. Additionally, it becomes apparent that wake interaction is necessary to derive the turbulence intensity when no other sensors are available. This observation conflicts with the objective of wake steering, in which wake interactions are minimized. Moreover, the analysis indicates that more complicated, unstructured wind farm layouts yield a higher observability than structured layouts due to the increase in unique wake interactions. Inferring from this theoretical analysis, it is apparent that sufficient observability of the ambient conditions in wind farms is not guaranteed. Therefore, estimation algorithms require careful implementation, taking this (lack of) observability into account.

Based on the gathered insights on observability, a closed-loop steady-state and model-based wind farm control solution is devised. As the first of its kind in the literature, this wind farm control algorithm is then tested in a high-fidelity simulation with time-varying inflow conditions. The wake steering algorithm leads to an averaged 1.4% increase in power production for a 6-turbine wind farm. The proposed controller proves to

be resilient to more realistic, time-varying inflow conditions and thereby solidifies itself as the first realistic, closed-loop wind farm control solution for yaw-based wake steering.

While the proposed closed-loop wind farm controller results in a net gain in the wind farm power production, there are time periods in which the power production was several percents lower than in locally greedy operation. These losses mainly occur when the inflow conditions change in time, leading to larger model mismatches and erroneous turbine misalignment.

To address the power production losses found during inflow changes, these temporal dynamics should be incorporated into the controller. Fundamentally, this implies that the focus should be shifted from steady-state towards dynamic surrogate wind farm models, which brings us to the third research question. The third research question is formulated as “*What benefits do dynamic surrogate models have for wind farm control and how can they be implemented?*” While beneficial to modeling accuracy, the main challenge of dynamic wind farm models is the significant increase in computational cost and algorithm complexity, which prevent its adoption in practice.

This thesis addresses one component of dynamic-model-based wind farm controller synthesis, being the development of a real-time model adaptation (state estimation) algorithm. The dynamic model employed is the nonlinear dynamic surrogate model WFSim and the estimator is based on an Ensemble Kalman filter. High-fidelity simulations show that the algorithm is able to accurately reconstruct the second-to-second two-dimensional flow field in a virtual 9-turbine wind farm. Moreover, the proposed algorithm is several orders of magnitude faster than the industry standard with a computational cost of 1.2 s per iteration, at a negligible loss in accuracy. Additionally, the estimator is tested assuming exclusively readily available measurements in the wind farm, which makes the algorithm feasible for practical implementation. This work provides an essential building block towards a reliable, practicable, closed-loop wind farm control solution with dynamic surrogate models.

In conclusion, this dissertation has matured the steady-state and model-based wind farm control concept on various fronts, essentially addressing the issue of model uncertainty. Closed-loop model-based wind farm control is paving the way for practicable wake steering. The contributions in this dissertation greatly advance the status quo of wind farm control solutions and their applicability in commercial wind farms.

SAMENVATTING

Klimaatverandering is een van de grootste actuele dreigingen voor de mensheid. In 2016 hebben 195 landen het *Accord de Paris* getekend, waarin ze beloven om de toename in de gemiddelde wereldwijde temperatuur te beperken tot 1.5°C ten opzichte van het pre-industriële tijdperk. De transitie van fossiele brandstoffen naar duurzame energiebronnen speelt hierin een essentiële rol. Windenergie wordt gezien als een van de meestbelovende duurzame energiebronnen en de voorspelling is dat het de grootste elektriciteitsbron van de Europese Unie zal zijn in 2050. Echter, economische en maatschappelijke redenen beperken momenteel de installatie van veel windturbines en windparken. Een manier om windenergie efficiënter te maken is door het gebruik van windparkregeling, waarin de regelaars van windturbines in een windpark met elkaar afgesteld wordt om zodanig een gezamenlijk doel te behalen. Windparkregelaars hebben de potentie om de energieopbrengst van windparken te verhogen, de mechanische belastingen in windturbines te verminderen, en windparken bruikbaar maken voor de stabilisatie van het elektriciteitsnetwerk.

Er bestaan verschillende soorten windparkregelaars, gebruikmakende van een verscheidendheid aan vrijheidsgraden en met uiteenlopende doelstellingen. De nadruk in deze dissertatie ligt op windparkregelaars die energieopbrengsten maximaliseren via het sturen van zoggen. Hierin wordt gebruik gemaakt van het feit dat zoggen verplaatst kunnen worden door stroomopwaartse windturbines een scheefstand te geven. Dit principe kan gebruikt worden om zoggen weg te sturen van stroomafwaartse turbines en zodoende de totale energieopbrengst te verhogen. In de afgelopen tien jaar is deze techniek gegroeid van een theoretisch idee tot een industrieel product. Het concept werd aanvankelijk getest met computermodellen, daarna via windtunnelexperimenten, en uiteindelijk in echte windparken. Terwijl twee van de drie veldexperimenten in de literatuur winsten laten zien met dit regelconcept, is het alleen getest voor twee-turbine interacties en specifieke atmosferische condities. Hierdoor blijft het onduidelijk in hoeverre zogsturing de totale jaarlijkse energieproductie van een commercieel windpark ten gunste brengt. Het derde veldexperiment beschreven in de literatuur omvat zogsturing gedurende een jaar op een commercieel windpark met zes windturbines. Alhoewel de auteurs winsten laten zien, blijkt uit hun analyse dat de toename in gemiddelde totale opbrengst verwaarloosbaar is. Derhalve is het succes van deze windparkregelaars onzeker. Deze dissertatie richt zich op deze wetenschappelijke kloof. De doelstelling van dit onderzoek is geformaliseerd als

“Het praktisch en robuust maken van windparkregelaars voor de maximalisatie van energieopbrengst via een model-gebaseerd en gesloten-lus framework”

Deze doelstelling is opgedeeld in drie onderzoeksvragen. De eerste onderzoeksvraag

is gedefinieerd als *“Wat zijn de limitaties van de huidige, open-lus windparkregelaars?”* Gevolglijk presenteert deze dissertatie de resultaten van een veldexperiment van zogsturing in een commercieel windpark aan de Italiaanse kust. In dit experiment worden twee- en drie-turbine interacties bestudeerd, waarin de eerste twee turbines onder een scheefstand gezet worden. De scheefstanden worden berekend door een simpel, tijds-onafhankelijk rekenmodel van het windpark, en veranderen afhankelijk van de windrichting, windsnelheid en turbulentie intensiteit. De resultaten laten een significante toename in de energieopbrengst zien voor bepaalde omstandigheden, met toenames in de energieopbrengst tot 35% voor twee-turbine interacties en tot 16% voor drie-turbine interacties. Echter, de resultaten laten ook verliezen zien door foutieve scheefstand. Deze foutieve scheefstand is verondersteld te komen door onnauwkeurigheden in het simpele rekenmodel van het windpark. Een grote bron van modelfouten zit in de gemiddelde relatie tussen the energieopbrengst van een wind turbine en zijn respectievelijke scheefstand. Ook maakt dit rekenmodel grote aannames over de complexe topologie en de effecten van turbines met verschillende hoogtes op de luchtstroom. Deze aannames zijn niet gevalideerd, en gaan daardoor gepaard met grote onzekerheden. Verder negeert dit rekenmodel belangrijke tijdsafhankelijke effecten zoals variaties in de wind en dynamica in de windturbines.

Met het oog op deze resultaten worden er twee additionele onderzoeksvragen gedefinieerd in deze dissertatie. De tweede onderzoeksvraag is geformuleerd als *“Kunnen wij robuustheid voor modelonzekerheid realiseren in moderne windparkregelaars via het actief terugkoppelen van metingen in de regelaar, en indien ja, hoe?”* De gemeten verliezen in het experiment in acht nemende, focust deze dissertatie zich op de transitie van open-lus naar gesloten-lus windparkregelaars. Hierin worden actief metingen gebruikt om de atmosferische condities af te schatten en het rekenmodel aan te passen zodat het accuraat de windparkinteracties blijft beschrijven.

Een nieuw theoretisch framework is geïntroduceerd in dit onderzoek die een kwantitatieve maatstaaf geeft aan de observeerbaarheid van een windpark. Deze maatstaaf duidt aan in hoeverre de atmosferische condities afgeleid kunnen worden van de beschikbare metingen. De analyse onthult welke metingen nodig zijn om de omgevingscondities accuraat te reconstrueren gebruikmakende van tijdsafhankelijke rekenmodellen. De resultaten bevestigen de opvatting dat lokale windsnelheids- en windrichtingsschattingen van windturbines nodig zijn om windparkgemiddelde schattingen te maken van de windsnelheid en windrichting. Verder wordt duidelijk dat zoginteractie nodig is om de turbulentie intensiteit af te schatten als er geen extra metingen beschikbaar zijn. Dit conflicteert met de doelstelling van windparkregeling, waarin zoginteractie geminimaliseerd wordt. Bovendien toont de analyse aan dat gecompliceerde, ongestructureerde windparktopologiën een hogere observeerbaarheid opleveren dan gestructureerde topologiën door de toename aan unieke zoginteracties. De theoretische analyse in deze dissertatie bewijst dat de mate van observeerbaarheid in windparken veel varieert, en dat de omgevingscondities voor veel situaties niet afgeschat kunnen worden. Daarom is het van belang dat schattingsalgoritmes dit (gebrek aan) observeerbaarheid

in acht nemen.

Met de inzichten van de observeerbaarheidsanalyse is er vervolgens een gesloten-lus windparkregelaar ontworpen, gebaseerd op een tijdsonafhankelijk rekenmodel van het windpark. Als de eerste in de literatuur is deze regelaar getest in een geavanceerd rekenmodel, onderworpen aan tijdsvariërende omgevingscondities. Het zogsturende algoritme leidt tot een toename van 1.4% in de energieopbrengst van een 6-turbine windpark. Deze regelaar bewijst zijn robuustheid voor realistische, tijdsvariërende windcondities, en vestigt zich hiermee als de eerste realistische, gesloten-lus regelaar voor zogsturing.

Terwijl het voorgestelde algoritme gemiddeld een toename in energieopbrengst oplevert, zijn er ook tijdsperiodes waarin er verliezen in de energieopbrengst gemeten worden. Deze verliezen worden verondersteld voort te komen uit tijdsveranderingen in de omgevingscondities, wat leidt tot grotere fouten in het simpele rekenmodel en tot foutieve scheefstanden op de turbines.

Om deze verliezen aan te kaarten, moet tijdsdynamica in acht genomen worden in het simpele rekenmodel van het windpark. Dit houdt in dat de focus van tijdsonafhankelijke naar tijdsafhankelijke rekenmodellen verschuift moet worden. De derde onderzoeksvraag is daarom geformuleerd als *“Welke voordelen hebben tijdsafhankelijke rekenmodellen voor windparkregeling, en hoe kunnen ze worden toegepast?”* Ondanks dat tijdsafhankelijke rekenmodellen beter het gedrag kunnen modelleren van windparken, komen ze gepaard met twee uitdagingen. De eerste uitdaging is dat tijdsafhankelijke modellen veel meer rekenkracht vereisen en daarom vaak langzaam zijn. De tweede uitdaging is de hoge complexiteit van het ontwerpen van regelaars die gebruik maken van zulke rekenmodellen. Deze uitdagingen hebben het gebruik van zulke modellen in de praktijk tot op heden belet.

De laatste contributie van deze dissertatie focust zich op een component van gesloten-lus regelaars met tijdsafhankelijke modellen, namelijk een algoritme die de omgevingscondities afschat en het rekenmodel actief corrigeert met metingen. Dit algoritme maakt gebruik van het rekenmodel WFSim en een Ensemble Kalman filter. Simulaties in een geavanceerd rekenmodel tonen aan dat het algoritme erin slaagt om het stromingsveld van een 9-turbine windpark af te schatten. Bovendien is het algoritme meerdere ordegroottes sneller dan de standaard algoritmes uit de literatuur met een rekentijd van 1.2 s per iteratie, zonder ten koste te gaan van de precisie. Verder is dit algoritme getest met metingen die standaard al beschikbaar zijn in windparken, wat de implementeerbaarheid bevordert. Dit werk verschaft een essentieel bouwblok voor het maken van realistische, bruikbare, gesloten-lus windparkregelaars met tijdsafhankelijke modellen.

In conclusie heeft deze dissertatie aantoonbare verbeteringen ontwikkeld voor de huidige windparkregelaars. Fundamenteel gezien heeft dit onderzoek de statistische onzekerheid in rekenmodellen aangekaart. Gesloten-lus en rekenmodel-gebaseerde regelaars bieden een solide basis voor robuuste en praktiseerbare algoritmes voor zogsturing in windparken. De contributies in deze dissertatie brengen significante vooruitgang in

SAMENVATTING

de status quo van windparkregelaars en hun bruikbaarheid in commerciële windparken.

CONTENTS

Acknowledgements	v
Summary	ix
Samenvatting	xiii
1 Introduction	1
1.1 The decay of our planet	3
1.2 The era of an energy revolution	4
1.3 Wind turbines, wind farms and the inevitable wake losses	5
1.4 Wind farm control and its objectives	6
1.4.1 Load mitigation	7
1.4.2 Electricity grid frequency regulation	8
1.4.3 Power maximization	8
1.5 How does wind farm control work?	8
1.5.1 Wake mitigation (axial induction control)	8
1.5.2 Wake displacement (wake redirection control)	11
1.6 The various wind farm controller architectures	13
1.6.1 Conventional, locally greedy control	13
1.6.2 Open-loop model-based control	14
1.6.3 Closed-loop model-free control	15
1.6.4 Closed-loop model-based control	16
1.7 Wind farm modeling for control applications	17
1.7.1 Steady-state surrogate models	17
1.7.2 Dynamic surrogate models	19
1.8 The goals and contributions of this thesis	19
1.8.1 The scientific gap	19
1.8.2 The goals of this thesis	20
1.8.3 Thesis outline	21
2 Testing open-loop wake steering at a commercial onshore wind farm	23
2.1 Introduction	25
2.2 Methodology	27
2.2.1 The wind farm	27
2.2.2 Experiment design	28
2.2.3 Data acquisition	29

CONTENTS

2.3	Controller synthesis	30
2.3.1	Estimation of the ambient conditions	31
2.3.2	Optimization of the turbine control setpoints	32
2.4	Data processing	33
2.5	Results & discussion	35
2.6	Conclusions.	39
3	What can we estimate with the measurements available?	41
3.1	Introduction	43
3.2	Surrogate model: FLORIS	46
3.3	Methodology: introducing a measure of observability	47
3.3.1	Cost function in estimation	48
3.3.2	A quantitative measure for observability.	51
3.4	A comprehensive observability analysis for 3 wind farm layouts	54
3.4.1	Estimating ϕ and U_∞ under perfect knowledge of I_∞	55
3.4.1.1	Redundancy in the cost function: power and wind speed estimates	56
3.4.1.2	Using exclusively wind speed estimator measurements ($\lambda_p = 0$, $\lambda_U = 1$, $\lambda_\phi = 0$)	57
3.4.1.3	Using local wind speed and wind direction estimates ($\lambda_p = 0$, $\lambda_U = 1$, $\lambda_\phi = 10$).	57
3.4.2	The full estimation problem: estimating ϕ , U_∞ and I_∞	59
3.5	Conclusions.	61
4	Steady-state-model-based wind farm control under time-varying inflow	63
4.1	Introduction	65
4.2	The Simulator for Wind Farm Applications	68
4.3	Surrogate model	68
4.3.1	Model definition	68
4.3.2	Model tuning prior to controller synthesis	69
4.3.3	Model validation.	71
4.3.4	Introducing a wake recovery factor.	72
4.4	Controller synthesis.	73
4.4.1	Real-time model adaptation	73
4.4.2	Real-time control setpoint optimization	75
4.4.3	An overview	76
4.5	Simulation results.	76
4.5.1	Model adaptation performance	76
4.5.2	Setpoint optimization performance	80
4.5.3	A deeper look into the yaw actuator duty cycle and structural loads	81
4.6	Conclusions.	83

5	Online model calibration for a dynamic surrogate model	85
5.1	Introduction	87
5.2	The surrogate model	89
5.3	Online model calibration	91
5.3.1	Challenges	92
5.3.2	General formulation	92
5.3.3	Extended Kalman filter (ExKF)	93
5.3.4	The Unscented Kalman filter (UKF)	94
5.3.5	The Ensemble Kalman filter (EnKF)	96
5.3.5.1	Challenges in the EnKF for small number of particles	97
5.3.6	Synthesizing an online model calibration solution	99
5.4	Results	100
5.4.1	SOWFA	100
5.4.2	2-turbine ALM with turbulent inflow	101
5.4.2.1	A comparison of the KF variants for state estimation	101
5.4.2.2	A comparison of sensor configurations	102
5.4.2.3	Joint state-parameter estimation	104
5.4.3	9-turbine ALM with turbulent inflow	105
5.5	Conclusions	108
6	Conclusions and recommendations	111
6.1	Conclusions	113
6.2	Recommendations	115
	References	121
A	Appendix: Additional field experiment look-up tables	137
B	Appendix: Yaw-power relationship for a GE 1.5s turbine	139
C	Appendix: Addressing irregularities in the degree of observability	141
	Curriculum Vitae	145
	List of Publications	147

1

INTRODUCTION

ABSTRACT

Climate change is one of the most imminent threats to society today. In 2016, 195 countries signed the Paris agreement, in which the signed parties commit to limiting the increase of the global average temperature to 1.5° C above pre-industrial levels. In its pursuit, wind energy is forecasted to become the European Union's single largest source of electricity by 2050. However, wind energy often remains infeasible compared to existing non-renewable energy sources for financial and societal reasons. One strategy to improve its feasibility is through wind farm control, in which the turbines inside a wind farm are controlled in a coordinated fashion to achieve a predefined objective. Concepts from this field show real potential in improving the performance of wind farms (e.g., annual energy yield or load mitigation). Some of the existing concepts have been tested in field experiments, showing both performance improvements and losses for a variety of situations. The losses are hypothesized to be due to inaccurate estimation of the ambient conditions on one hand and the usage of erroneous mathematical models in the wind farm control algorithms on the other hand. This chapter presents how this thesis will address these scientific gaps. More formally, the overarching objective of this thesis is "Maturing wind farm control technologies for power maximization in a model-based closed-loop framework towards real-world practical applicability." The research gaps and thesis goals are presented in more detail at the end of this chapter, preceded by a general introduction to wind farms and wind farm control.

The author has published various works of literature that provide an introduction to the topic of wind farm control. This chapter combines essential elements of these works of literature to provide a coherent introduction to the topic of this dissertation.

- S Boersma, B M Doekemeijer, P M O Gebraad, P A Fleming, J R Annoni, A K Scholbrock, J A Frederik, and J W van Wingerden. A tutorial on control-oriented modeling and control of wind farms. In *American Control Conference (ACC)*, pages 1–18, 2017. URL <http://doi.org/10.23919/ACC.2017.7962923>
- B M Doekemeijer, P A Fleming, and J W van Wingerden. A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model. In *American Control Conference (ACC)*, pages 2825–2836, 2019b. URL <https://doi.org/10.23919/ACC.2019.8815126>
- B M Doekemeijer, S Boersma, J R Annoni, P A Fleming, and J W van Wingerden. Wind plant controller design. In Paul Veers, editor, *Wind Energy Modeling and Simulation; Volume 1: Atmosphere and plant*, Energy Engineering, chapter 7. The Institution of Engineering and Technology (IET), 2019a. ISBN 978-1-78561-521-4

1.1. The decay of our planet

The 29th of July in 2019 was depicted as the *Earth Overshoot Day* by the Global Footprint Network, being the day that humanity consumed all of the Earth's natural resources for the year if we were to use the planet in a sustainable manner. The fact that Earth Overshoot Day is in late July equates to the global consumption of about 1.75 Earths each year. The natural resources and their consumption are unevenly distributed among countries, with The Netherlands using 6 times the sustainable level of its resources. This number is 10 for Luxembourg, 19 for Israel and a staggering 101 for Singapore (Global Footprint Network, 2019). The Earth Overshoot Day comes earlier than ever before and has been growing since the one-Earth consumption threshold was crossed in 1970. Figure 1.1 shows the trend of the Earth Overshoot Day from 1965 to 2019.

The Earth Overshoot Day is characterized by five features: cities, energy, food, population and our planet's biological resources. Accordingly, cities must be designed to accommodate the population and enable natural resources to thrive; our energy must be harvested from renewable, carbon-neutral energy sources; the production of food and clean drinking water must meet the growing population's needs in a sustainable manner; and the human population at large must be distributed and facilitated such that everyone has access to clean drinking water, shelter and healthy nutrition.

While the *Earth Overshoot Day* metric has been criticized in the literature (e.g. Blomqvist et al., 2013; Giampietro and Saltelli, 2014), the consensus remains that our current lifestyle cannot be sustained on this planet. Specifically, greenhouse gas emissions must be reduced for a durable future.

A 2014 report by the Intergovernmental Panel on Climate Change (IPCC) – a United-Nations-governed panel that assesses the science related to climate change – concludes:

“Human influence on the climate system is clear and recent anthropogenic emissions of greenhouse gases are the highest in history. Recent climate changes have had widespread impacts on human and natural systems.”

– Pachauri et al. (2014)

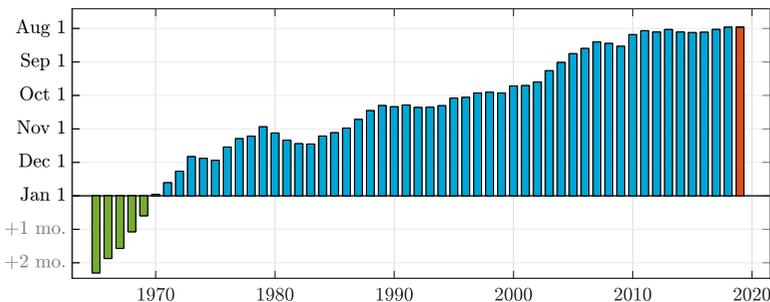


Figure 1.1: The annual date of *Earth Overshoot Day* over the period of 1965 to 2019. Human presence has been leaving an increasingly negative footprint on the planet from 1970 onwards. The data used for this figure has been retrieved from the Global Footprint Network (2019).

As human civilization continues the relentless emission of greenhouse gases, the habitability of our planet will further decay. Among others, a substantial number of animal species is now threatened with extinction, sea levels will rise, droughts may become more widespread and the air quality worsens. In turn, this may lead to a shortage in potable water and food, forced migration due to floods (among which are in The Netherlands, The Philippines, Brazil, India, the UK and China (Yeung, 2019)) and increased poverty (Masson-Delmotte et al., 2019).

1.2. The era of an energy revolution

In view of such news, the future for our society does not look bright. Fortunately, change is on the horizon. Over the past years, legislation has been adapting to tackle the problem of climate change. In 2016, 195 countries signed the *Accord de Paris*, the Paris Agreement (United Nations, 2016), in which the signed parties commit to pursuing the cease of the threat of climate change, migrate towards sustainable development and eliminate poverty by, among others,

“[...] holding the increase in the global average temperature to well below 2° C above pre-industrial levels and pursuing efforts to limit the temperature increase to 1.5° C above pre-industrial levels, recognizing that this would significantly reduce the risks and impacts of climate change.”

– United Nations (2016)

In its pursuit, legislation in Europe is introduced to mitigate pollution in a multitude of sectors, including the agricultural sector (e.g., European Commission, 2017), the transportation sector (e.g., Philbin and Raillant-Clark, 2017), the industrial sector (e.g., European Environment Agency, 2019) and the electricity sector (e.g., European Commission, 2019). Specifically, the European Commission has adopted one of the most aspiring policies on renewable energy in 2019, committing to a minimum of 32% of the produced energy to be from renewable sources by the year 2030 (European Commission, 2019). This renewable energy comes from a variety of sources, among which are wind, solar, hydro, tidal, geothermal and biomass.

The current and forecasted gross electricity production by energy source across the 28 European Union member states is shown in Figure 1.2. Wind energy is expected to become the single largest electricity source of the European Union by 2050. In fact, the second largest renewable energy source by 2050, solar, is hypothesized to be a factor 2 to 3 smaller. Accordingly, the Dutch government is doubling down on offshore wind energy with 1502 MW of wind energy capacity planned for the upcoming Borssele wind farms (currently being installed), a capacity of 3540 MW scheduled for the Hollandse Kust wind farms to be installed by 2021, 700 MW planned for the wind farm north of the Wadden Islands by 2022 and another 4000 MW of capacity planned for the IJmuiden wind farms by 2026 (Rijksdienst voor Ondernemend Nederland, 2019). In addition to the existing offshore wind farms, this sums up to 11.5 GW of wind energy by 2026. Under full operation, the Dutch offshore wind farms would be able to provide the equivalent

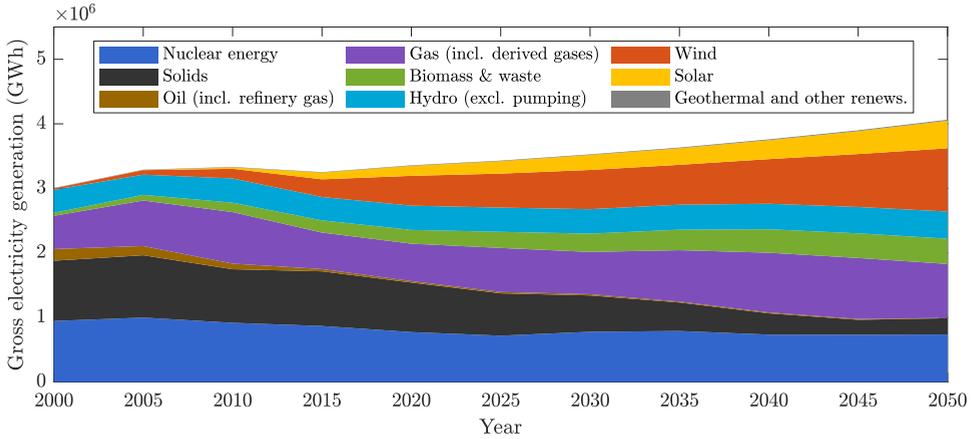


Figure 1.2: Gross electricity production in the EU-28 by energy source, from 2000 forecasted until 2050. Wind energy is expected to become the single largest source of electricity for the European Union by 2050. (European Commission, 2016)

amount of electricity consumed by 29 million households.¹

1.3. Wind turbines, wind farms and the inevitable wake losses

Modern, commercial wind turbines are designed to extract as much energy from a wind flow as physically possible. The theoretical upper limit for the energy extraction is described by Betz limit², which states that (van Kuik, 2007):

“[...] no more than 59% of the kinetic energy contained in a stream tube having the same cross-section as the disc area can be converted to useful work by the disc.”
– van Kuik (2007)

Basically, this means that a conventional, horizontal-axis wind turbine cannot extract more than 59% of the kinetic energy contained in a wind stream. This energy is converted to electrical energy through a generator. Interest in the improvement of wind turbine operation and efficiency remains widespread in the literature (e.g., Mulders, 2020), yet the largest developments in this field are made by the industry nowadays. Modern horizontal-axis wind turbines now approach the Betz limit, extracting in the order of 45 – 50% of the kinetic energy.

Over the past decade, the academic world has shifted its focus from wind turbines towards wind farms. While individual horizontal-axis wind turbines attain a high energy

¹Here, each household is assumed to consume 400 W around-the-clock (Energieleveranciers, 2019).

²Actually, the *Betz limit* should be called the *Lanchester-Betz-Joukowski limit*, as all three scientists independently discovered this theoretical limit around the same time in the early 20th century (van Kuik, 2007). However, the limit is most commonly referred to as *Betz limit* in academia and therefore continues to be named as such in this thesis.



Figure 1.3: A photograph by C. Steiness of the Horns Rev wind farm off the coast of Denmark in foggy conditions (Hasager et al., 2013). These unique atmospheric conditions allow us to see wake formation and wake interaction that are otherwise invisible under normal weather conditions.

conversion efficiency, the same cannot be said for wind farms. This is due to the formation of wind wakes, which are slower, more turbulent streams of air that emerge behind a turbine rotor during operation. These wakes often persist for several kilometers, causing efficiency losses on downstream turbines.

A photo of wake formation in the Horns Rev offshore wind farm is shown in Figure 1.3. Sørensen et al. (2006) estimated that wake formation in this commonly-spaced wind farm leads to an annual loss in the total energy yield of 12% compared to the hypothetical situation without wake formation. Similarly, Dahlberg and Thor (2009) estimated that wakes lead to a loss of 23% for the closely-spaced Lillgrund offshore wind farm. To further promote the financial competitiveness of wind as a renewable energy source, these wake losses must be addressed. Most commonly, this is done through *wind farm control*.

1.4. Wind farm control and its objectives

Wind farm control is a hypernym (*umbrella term*) for operational strategies in which the individual turbine's control settings are scheduled according to the behavior of other turbines within that wind farm. This is in contrast to the traditional way of operating wind turbines, *greedy control*, in which the influence of each wind turbine on neighboring units is neglected. In essence, wind farm control methods operate turbines in a coordinated manner to achieve a common objective.

At large, wind farm control works towards the minimization of the levelized costs of

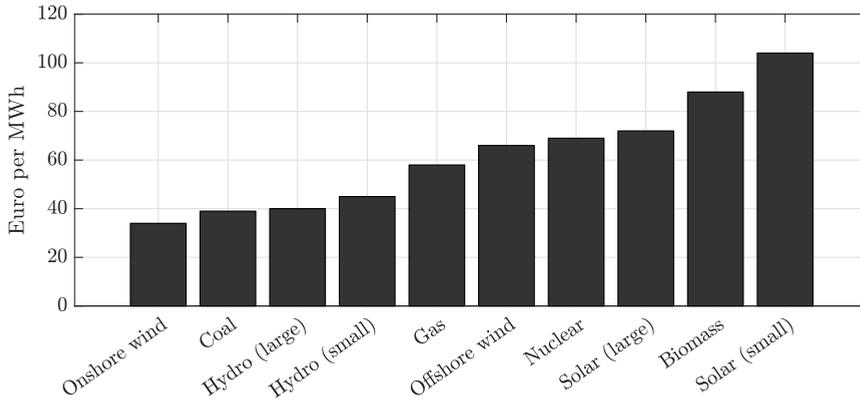


Figure 1.4: Levelized cost of energy by source for Germany in 2015 (Hansen, 2019). This figure shows that onshore wind is the least expensive form of energy wherever it is available. However, due to adverse effects of wind turbines such as shadow flicker, noise and visual pollution, social acceptance of onshore wind energy has been limited. Offshore wind remains to be almost twice as expensive as onshore wind energy and more expensive than electricity from coal and gas.

energy (LCOE). Fundamentally, the LCOE is a measure that allows cost-of-energy comparisons between various energy sources. In practice, the minimization of the LCOE implies the maximization of the energy yield and the minimization of the manufacturing, installation, maintenance and decommissioning costs of the wind energy system.

Figure 1.4 displays the LCOE for various energy sources in Germany, the largest country in the European Union. This figure shows the increasing competitiveness of wind energy, with onshore wind being the most cost-effective source of energy for Germany. However, adverse effects of wind turbines such as shadow flicker, noise and visual pollution limit the societal acceptance of wind energy onshore (Hickman, 2012). In effect, there is a trend towards the installation of wind farms at sea. Though, offshore wind energy is about twice as expensive as onshore wind energy and to date more expensive than the pollution-heavy, non-renewable energy sources coal and gas.

Wind farm control has the potential to further reduce the LCOE of wind energy. It does so through three subgoals, being the minimization of wear on the turbine structure, the integration with the electricity grid and the maximization of the annual energy yield. Each aspect is further elaborated upon in Sections 1.4.1 to 1.4.3. Subsequently, Section 1.5 introduces the two fundamental control methods used to attain these subgoals.

1.4.1. Load mitigation

The first goal of wind farm control is structural load mitigation. For example, one can leverage the concepts of wind farm control to distribute the structural loads between turbines in a farm. In this manner, turbines that have experienced more wear than others can be relieved, such that the turbine structures in a farm degrade at an equivalent rate. Such techniques may reduce turbine maintenance costs (Obdam et al., 2007), allow for

a lighter and less expensive design and increase turbine lifetimes (Kanev et al., 2018).

1.4.2. Electricity grid frequency regulation

The second goal of wind farm control is the integration of electricity from wind turbines with the electricity grid, commonly referred to as *active power control* or the provision of *ancillary grid services*. The electricity consumed in the electricity grid must match the electricity supplied to the grid. If the mismatch between generation and consumption is too large, black-outs and grid-connected machine failures may occur. Grid stability is currently provided by conventional sources such as hydraulic and thermal power plants which can quickly derate and uprate their power supply, while wind and solar energy sources were mostly exempt from this functionality (Díaz-González et al., 2014).

However, as the penetration of wind and solar energy in the electricity grid increases, there is a strong need for these renewable technologies to provide grid frequency regulation (Aho et al., 2012). The concepts of wind farm control has shown promising results in enabling wind farms to provide ancillary grid services (Ela et al., 2014). Moreover, this concept of active power control can straightforwardly be combined with a secondary goal such as reserve power maximization (e.g., Boersma et al., 2019b) or structural load minimization (e.g., Baros and Annaswamy, 2019; Kazda and Cutululis, 2019). This type of wind farm control has been particularly of interest for the industry.

1.4.3. Power maximization

The third and most common goal of wind farm control in academia is power maximization. The concept of wind farm control for power maximization has been demonstrated successfully in high-fidelity simulations (e.g., Gebraad et al., 2016), wind tunnel experiments (e.g., Campagnolo et al., 2016b) and field experiments (e.g., Doekemeijer et al., 2020a; Fleming et al., 2017b, 2019; Howland et al., 2019). For particular wake-loss heavy situations (i.e., particular wind directions and low turbulence), wind farm control has the potential to increase the energy yield by up to 20%. While the concept is yet to be demonstrated successfully throughout annual wind farm operation in the literature, the industry is starting to explore the commercialization of the concept (e.g., Siemens-Gamesa Renewable Energy, 2019). For these reasons, the primary focus in this thesis is on wind farm control for power maximization.

1.5. How does wind farm control work?

As outlined in Section 1.4, wind farm control has three main goals: load minimization, electricity grid integration and power maximization. On a technological level, these three objectives are realized by two distinct wind farm control methods: wake mitigation and wake displacement. Both approaches are described next.

1.5.1. Wake mitigation (axial induction control)

The first wind farm control method is wake mitigation, or commonly referred to as *axial induction control*. In this method, turbines are controlled to intentionally capture less energy by pitching the blades and changing the generator torque, as these actions reduce

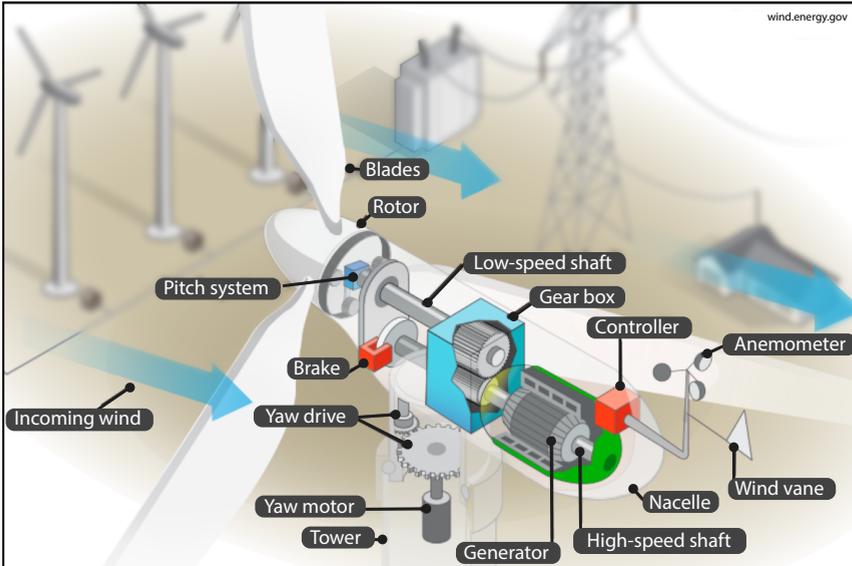


Figure 1.5: The insides of a horizontal axis wind turbine. Wind turbine and farm control algorithms rely on three degrees of freedom in the turbine: the pitch angle of each blade, the yaw angle and the generator torque. This image has been taken and modified from the U.S. Department of Energy.

the aerodynamic efficiency of the rotor. Figure 1.5 depicts the location of these control variables in the wind turbine structure, including an overview of the main components inside the wind turbine nacelle. The concept of axial induction control is visualized in Figure 1.6.

In the wind farm control literature, steady-state axial induction control has been explored extensively. Nonetheless, results in the literature (Annoni et al., 2016a; Campagnolo et al., 2016a) are inconclusive concerning its potential for power maximization. Instead, steady-state axial induction control is most often used for active power control (e.g., van Wingerden et al., 2017) and load mitigation (e.g., Kanev et al., 2018).

In more recent work (Frederik et al., 2020c; Munters and Meyers, 2018), the concept of dynamic (axial) induction control is explored, in which an upstream turbine follows a sinusoidal derating setpoint. This control concept, illustrated in Figure 1.7, leads to alternating pockets of slower and faster air to flow through the rotor plane. This approach has shown to induce additional wake recovery downstream, leading to a gain in the combined energy yield of a two-turbine array of up to 5%. On the other hand, this control strategy leads to a significant increase in the mechanical fatigue loads and the pitch actuator duty cycle on the upstream turbine due to the intentional periodic derating.

Moreover, Frederik et al. (2020a) demonstrate the novel concept of dynamical individual pitch control, which is an induction control method in which a directional moment is generated on the rotor plane by pitching the blades. Similarly, the flow applies

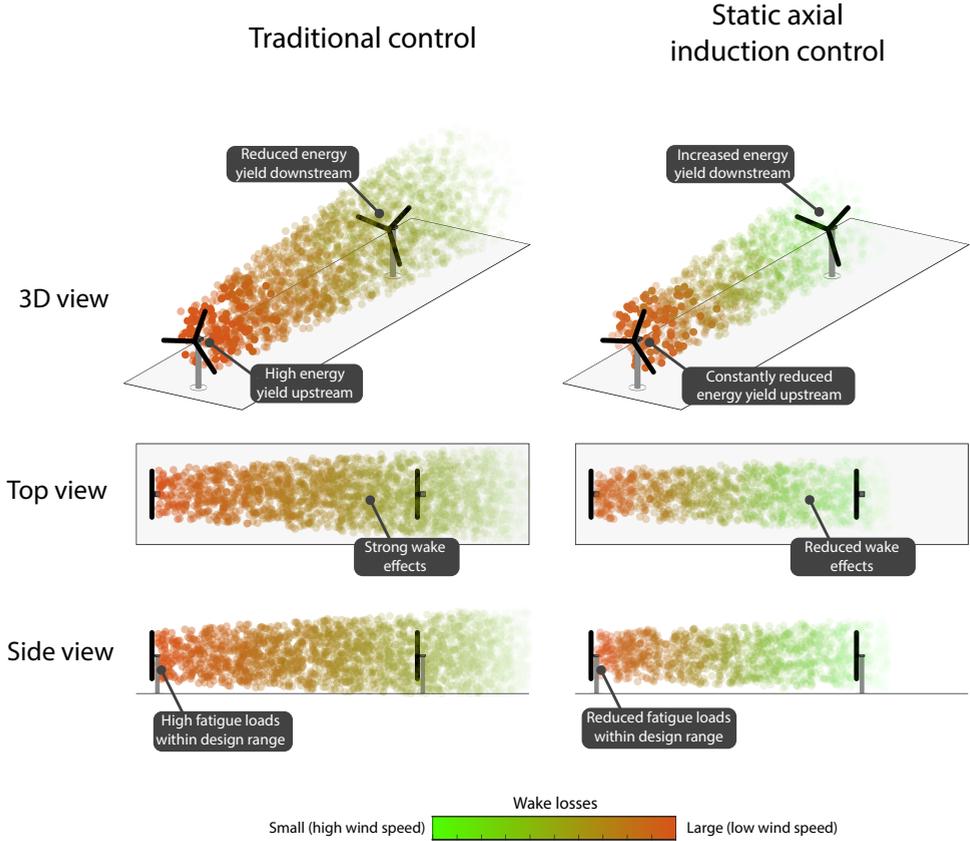


Figure 1.6: Static axial induction control demonstrated. On the left, the situation for traditional, locally greedy control is depicted. As the upstream turbine attempts to extract as much energy as possible, strong wakes trail behind the rotor and structural loads on the turbine are large. This leads to significant losses in the energy yield at the downstream turbine. By derating an upstream turbine (static axial induction control, depicted on the right), the structural loads reduce, wake effects diminish and the downstream turbine can achieve a higher energy yield. For clarity, the wake of the second turbine is not shown.

an opposite moment on the rotor. The direction of this moment slowly turns across the rotor plane, leading to a helical wake in the three-dimensional space. In simulation, this novel concept shows a notable increase in wake recovery and thereby power production at a downstream turbine. At the same time, the power and load fluctuations experienced by the upstream turbine are reduced compared to dynamic induction control. Though, the impact of such control strategies on the structural loads and blade pitch actuators are yet to be investigated in detail.

Static axial induction control leverages the common degrees of freedom of the turbine and does not operate outside of its design specifications. Therefore, the adoption of static axial induction control has not met much resistance from the industry and is cur-

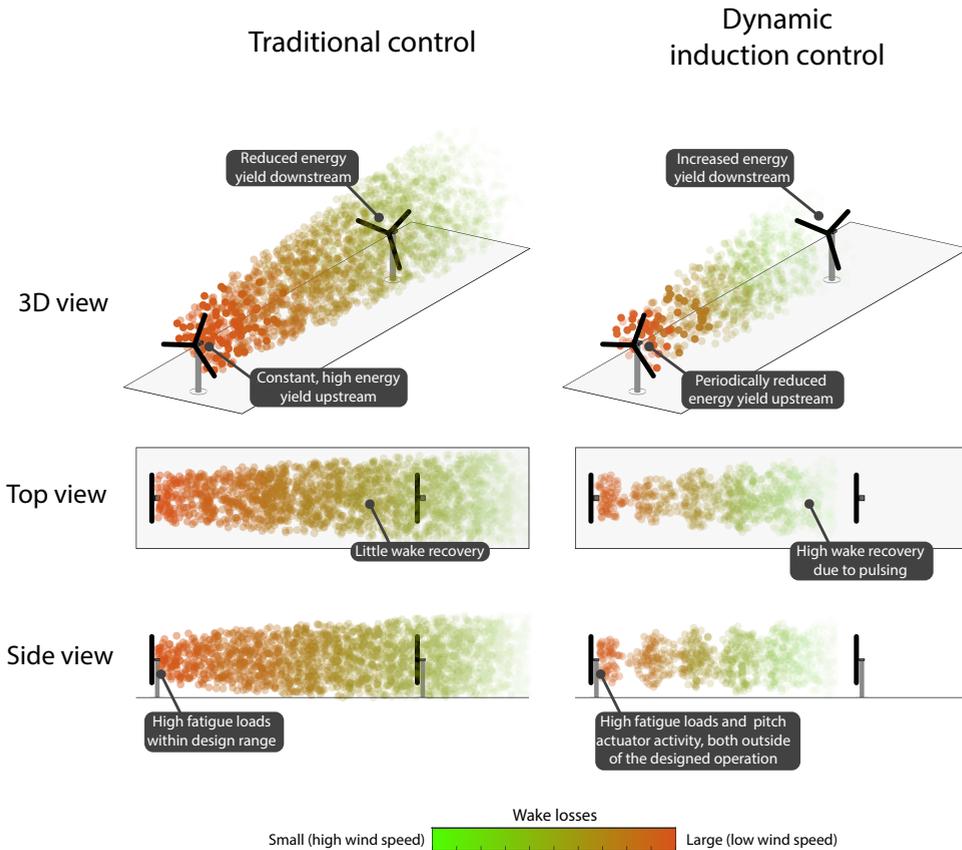


Figure 1.7: Dynamic (axial) induction control demonstrated. On the left, the situation for traditional, locally greedy control is depicted. As the upstream turbine attempts to extract as much energy as possible, strong wakes trail behind the rotor and structural loads on the turbine are large. This leads to significant losses in the energy yield at the downstream turbine. By constantly varying the amount of energy extracted by an upstream turbine (dynamic induction control, depicted on the right), additional wake recovery is induced downstream. This leads to a significant increase in the pitch actuator duty cycle and the fatigue loading on the upstream turbine, for which the turbine was not designed. However, wake effects diminish further downstream and the downstream turbine can achieve a higher energy yield. For clarity, the wake of the second turbine is not shown.

rently already employed in some wind farms for the provision of ancillary grid services.

1.5.2. Wake displacement (wake redirection control)

The second control method is wake displacement, or commonly referred to as *wake redirection control*. In this method, turbines are controlled to purposely capture less energy by exerting an uneven loading on the incoming airflow. This can be done by yawing the turbine away from the dominant wind direction (Jiménez et al., 2010), by tilting the rotor plane (Annoni et al., 2017), or by pitching the turbine blades (Fleming et al., 2014). The

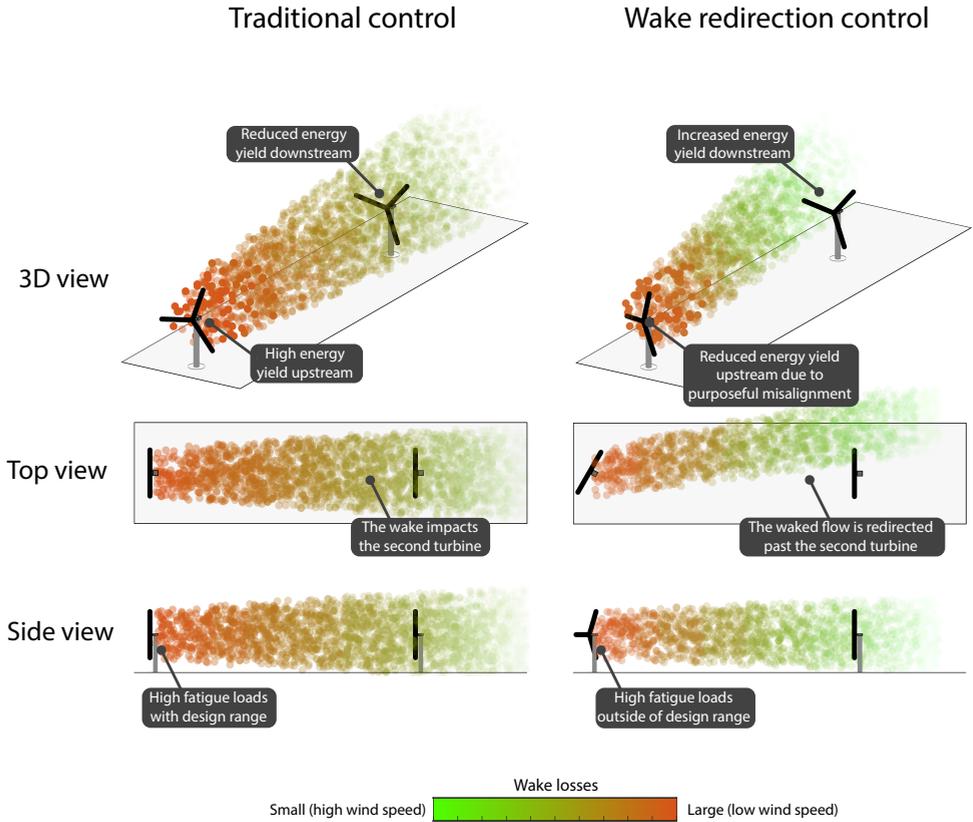


Figure 1.8: Yaw-based wake redirection control demonstrated. On the left, the situation for traditional, locally greedy control is depicted. As the upstream turbine attempts to extract as much energy as possible, strong wakes trail behind the rotor and structural loads on the turbine are large. This leads to significant losses in the energy yield at the downstream turbine. By purposely yawing the rotor plane away from the mean wind direction (yaw-based wake steering, depicted on the right), the wake can be steered away from a downstream turbine. This leads to an increase in the yaw actuator duty cycle and may lead to an increase in the fatigue loading on the upstream turbine, for which the turbine was not designed. However, wake effects are displaced from a downstream turbine and the downstream turbine can thereby achieve a higher energy yield. For clarity, the wake of the second turbine is not shown.

concept of rotor tilt for wake steering has not been explored in much detail as it would require a complete redesign of the wind turbine. Additionally, the concept of wake steering by blade pitch has not shown much promise in the literature (Fleming et al., 2014) and this degree of freedom is typically used for load alleviation instead (Bossanyi, 2003; Mulders and van Wingerden, 2019). The most common method for wake displacement is yaw-based wake redirection control and is visualized in Figure 1.8.

Yaw-based wake steering is most commonly used for power maximization (Boersma

et al., 2017). Due to its inherently slow actuation and large load on the yaw motors, it has no application in active power control. Additionally, while some studies suggest that turbine loads may decrease with particular yaw misalignments (Reyes et al., 2019), the literature has not been conclusive on this topic. Initial work on yaw-based wake steering for power maximization has shown its potential in simulation (Fleming et al., 2014; Jiménez et al., 2010) and more recently in wind tunnel experiments (Bastankhah and Porté-Agel, 2016; Campagnolo et al., 2016a,b) and field experiments (Fleming et al., 2017b, 2019; Howland et al., 2019).

For particular wind directions and low turbulence levels, wake-loss-heavy situations arise, in which wake steering has shown the ability to increase the collective power production of that wind farm by up to 15%. As wind farms are already designed with wake losses in mind, yaw-based wake steering is projected to increase the annual energy production of a wind farm by up to 1% in practice. While this number might seem insignificant, it is a notable improvement in wind farm performance at a low investment cost. The method is now slowly being adopted by the industry, aiming to be sold to wind farm operators as a service (e.g., Siemens-Gamesa Renewable Energy, 2019). Because of its high potential, the focus in this thesis is on yaw-based wake steering for power maximization.

1.6. The various wind farm controller architectures

With the challenges, objectives and methods presented in Sections 1.3 to 1.5, the focus can now be shifted towards the practical synthesis of wind farm control solutions. Note that most commercial wind farms facilitate ancillary grid services, addressed previously in Section 1.4.2, and therefore have a wind farm control framework implemented (for more information, see Sørensen et al., 2005). However, the focus in the remainder of this dissertation is on yaw-based wake steering for power maximization, and in this context, a control framework is absent in commercial wind farms. The remainder of the dissertation must be read with this in mind.

While the difference is sometimes somewhat unclear, wind farm control algorithms can roughly be distinguished into *model-based* and *model-free* solutions. Model-based solutions leverage a mathematical model of the wind farm in real-time to decide the next control action. Model-free methods abstractify the wind farm as a black box and aim to converge to an optimal control action by looking at the measured system responses directly. Additionally, architectures may be separated in *closed-loop* (feedback) and *open-loop* (feedforward) wind farm control solutions. Feedback algorithms use wind turbine measurements in real-time for the determination of the control policy. Feedforward algorithms exclusively rely on prior information. Each architecture is briefly discussed next.

1.6.1. Conventional, locally greedy control

The first architecture is that of the industry-standard, locally greedy wind farm control, depicted in Figure 1.9. Here, *locally greedy* refers to the fact that each turbine is controlled on an individual level to maximize its own energy yield and minimize fatigue

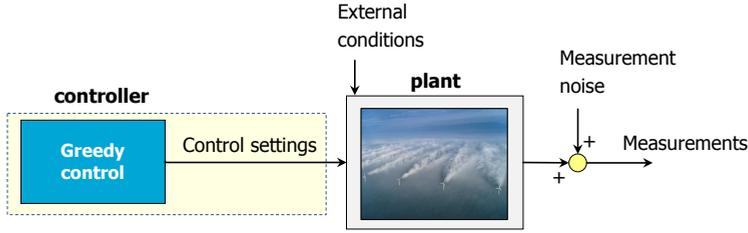


Figure 1.9: The architecture for locally greedy wind farm control. Here, turbines are controlled on an individual level, ignoring their impact on surrounding turbines.

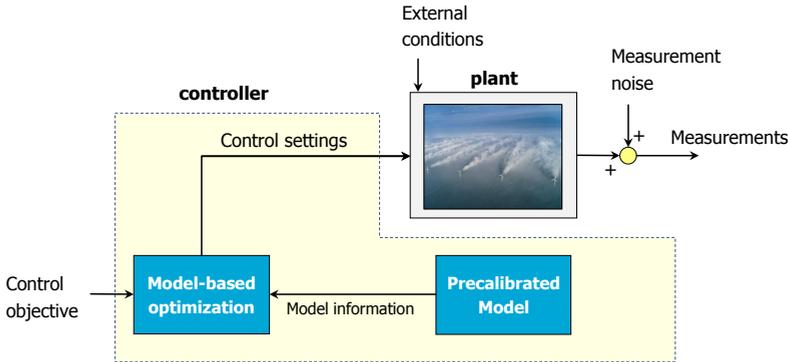


Figure 1.10: The architecture for open-loop model-based wind farm control. Here, turbines are controlled according to a simplified mathematical model of the wind farm. The effects of the applied control policy is not measured or taken into account in deciding the next control policy and therefore the method is highly sensitive to the accuracy of the surrogate wind farm model.

loads, ignoring the impact on other turbines through wake formation. For a more in-depth introduction to traditional wind turbine control, the reader is referred to Pao and Johnson (2009).

1.6.2. Open-loop model-based control

The second architecture encompasses open-loop model-based wind farm controllers, as depicted in Figure 1.10. These control methods use a (possibly precalibrated) surrogate model of the wind farm to determine the turbine control settings. These surrogate wind farm models capture the effect of certain turbine control settings (e.g., blade pitch or nacelle yaw) on performance parameters (e.g., power production or structural loads). These surrogate models trade off accuracy with computational cost to enable real-time optimization.

Most of the yaw-based wake steering algorithms in the literature are of such architecture (Campagnolo et al., 2016a; Fleming et al., 2017b, 2019; Gebraad et al., 2016; Howland et al., 2019). These control algorithms assign control settings to each turbine according

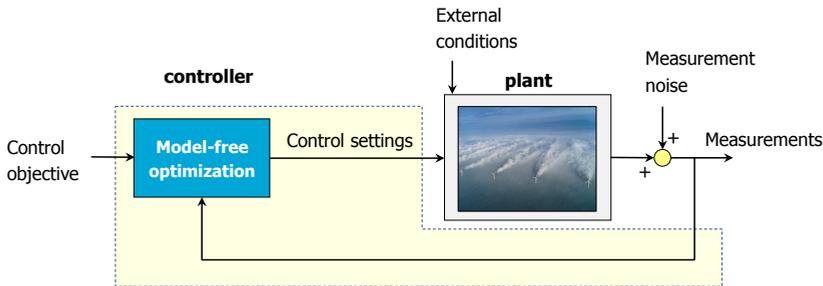


Figure 1.11: The architecture for closed-loop model-free wind farm control. Here, the wind farm is considered a black box system with inputs and outputs. The controller directly manipulates the system inputs to achieve a certain objective, measured according to the system outputs. While this architecture surpasses the issue of establishing an accurate surrogate wind farm model, these methods face challenges with convergence because of constantly changing inflow conditions and large time delays induced by wake formation.

to the measured wind direction and wind speed, but do not actually evaluate the impact of the applied control settings on the turbine's power production. Therefore, despite the fact that turbine measurements are used in deciding the wind farm control policy, it is considered an open-loop wind farm control solution. Since these methods do not evaluate the success of the applied control settings, they are highly sensitive to the accuracy of the surrogate wind farm model. Erroneous or poorly calibrated wind farm models may lead to poor wind farm performance in practice, even when compared to traditional, locally greedy control (Annoni et al., 2016a; Doekemeijer et al., 2019b).

1.6.3. Closed-loop model-free control

The third architecture comprises closed-loop model-free wind farm controllers, depicted in Figure 1.11. Fundamentally, these controllers consider the wind farm as a black box system with inputs and outputs, upon which the inputs are to be manipulated to achieve a certain objective according to the system outputs. Model-free methods circumvent the issue of capturing the complicated temporal and spatial dynamics of wind farms in a simplified mathematical model by directly optimizing the measured outputs of the system. This concept has been successfully demonstrated for power maximization through yaw-based wake steering in a wind tunnel subject to static inflow conditions in Campagnolo et al. (2016b). Furthermore, the concept has been demonstrated for power maximization through axial induction control by Ciri et al. (2017a, 2019). Though, this work is more towards wind turbine control than wind farm control as wake interactions are neglected.

However, there is a drawback to model-free wind farm control. As the freestream wind direction and wind speed in the wind farm continuously change, so do the optimal yaw misalignment angles. Moreover, it often takes several minutes before the effect of a change in the yaw angle is seen on a downstream turbine due to time delays from

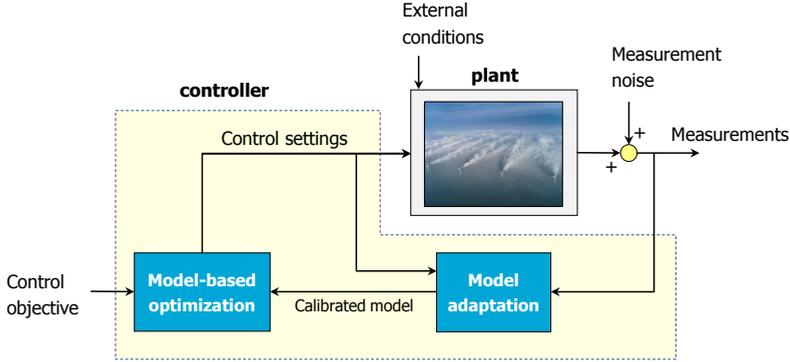


Figure 1.12: The architecture for closed-loop model-based wind farm control. This architecture combines the benefits of closed-loop model-free with open-loop model-based control, being quick convergence rates and resilience to model uncertainties, respectively. Here, turbines are controlled according to a simplified mathematical model of the wind farm. Additionally, the effects of the applied control policy is measured and taken into account in deciding the next control policy, either directly or through recalibration of the surrogate wind farm model (exemplified here).

wake propagation. Given the fact that the inflow is continuously changing, model-free optimization methods face challenges with convergence and often appear too slow for real-time power maximization, e.g., through yaw-based wake redirection (Boersma et al., 2017).

Instead, model-free wind farm control methods are often used for active power control. Namely, wake effects need not be accounted for explicitly in a wind farm control algorithm where power signal reference tracking is the objective (e.g., Boersma et al., 2019b; van Wingerden et al., 2017). Additionally, the generator torque is quickly updated or downrated to follow reference signals that vary much faster than timescales of wake propagation. Here, wakes are often considered to be a disturbance rather than an effect to be included in the controller.

1.6.4. Closed-loop model-based control

The fourth architecture encircles closed-loop model-based wind farm controllers and is visualized in Figure 1.12. These control solutions combine the benefits of open-loop model-based control, being the computational efficiency and quick convergence, with the benefits of closed-loop model-free control, being resilience to modeling uncertainties.

As these control solutions are typically more complex than open-loop and model-free solutions, literature on this topic remains somewhat sparse. In theory, this architecture should outperform the previously discussed architectures. However, its main challenges are controller stability and the computational cost involved. Controllers of this architecture have been employed for active power control (e.g., Bay et al., 2018; Boersma et al., 2018c; Kazda and Cutululis, 2019; Shapiro et al., 2018; Spudić et al., 2015), load mitigation (e.g., Boersma et al., 2018c; Spudić et al., 2015) and power maximization (e.g.,

Gebraad et al., 2015; Raach et al., 2019). Because of the lack of research in this field, this thesis aims to further mature this controller architecture to unleash the full potential of closed-loop model-based wind farm control. The thesis goals are further formalized in Section 1.8, after a brief discussion on surrogate wind farm models in Section 1.7.

1.7. Wind farm modeling for control applications

The performance of model-based wind farm control solutions is closely linked to the accuracy and computational cost of the applied mathematical model. According to these measures, models vary from low to high fidelity. Low-fidelity models have a lower accuracy, a low computational cost and are typically used for controller synthesis and wind farm analysis (e.g., AEP predictions). High-fidelity models have a higher accuracy, a high computational cost and are commonly used for offline applications such as wind farm flow and turbine simulations. While Goit and Meyers (2015); Munters and Meyers (2017) present a wind farm control solution using a high-fidelity model, these authors also state that

“[...] Finally, the current optimal control approach allows one to benchmark control potential, but is not practicable for use as a real-time controller.”

– Goit and Meyers (2015)

In other words, the computational cost involved would be much too high for real-time optimization. No other research groups have directly applied high-fidelity models for model-based wind farm control. Instead, practical model-based control solutions leverage low-fidelity models with an affordable computational complexity. These low-fidelity surrogate wind farm models can be separated into static and dynamic models. Both are briefly addressed in this section.

1.7.1. Steady-state surrogate models

The main focus in the literature on surrogate modeling for wind farms has been on steady-state models. In these models, all temporal dynamics are neglected (or rather, a temporal average is captured) and the model is effectively an input-output mapping. An example of the flow field of a popular steady-state model is shown in Figure 1.13. In these models, the inputs may be the turbine control settings and the freestream wind field and the outputs may be the turbine power signals and the developed flow field, for example. This allows for a low computational cost and thereby makes the approach suitable for real-time control. The most popular application of steady-state surrogate wind farm models lies in power maximization through yaw-based wake steering (e.g., Doekemeijer et al., 2019b; Fleming et al., 2017b, 2019; Gebraad et al., 2016; Howland et al., 2019) and axial induction control (e.g., Annoni et al., 2016a; Santoni et al., 2015). See Tables 2 to 5 in Kheirabadi and Nagamune (2019) for an overview of the studies performed on this topic.

In addition to power maximization, steady-state surrogate wind farm models are also used for load mitigation (e.g., van Dijk et al., 2017), though to a significantly lower de-

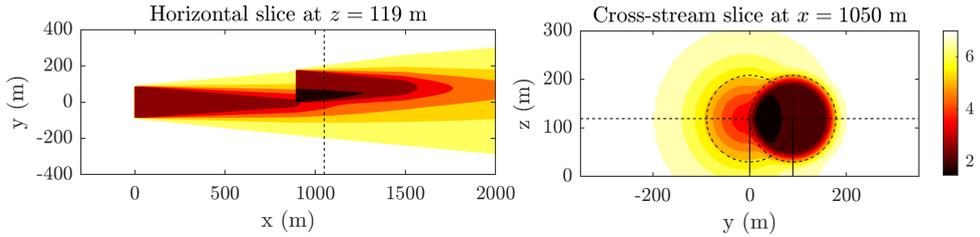


Figure 1.13: Three-dimensional flow field for a wind farm with two DTU 10MW wind turbines as predicted by FLORIS. The figure shows a horizontal slice of the flowfield at hub height and a cross-stream slice at some distance downstream. The Gaussian nature of the wakes is clearly seen.

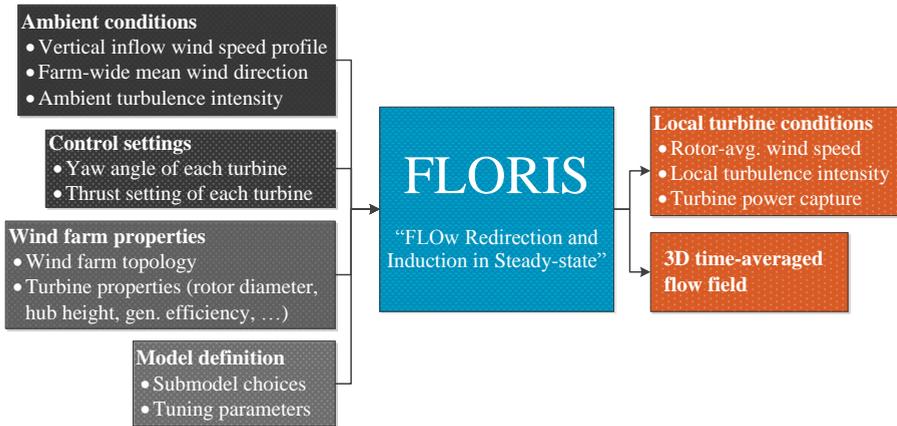


Figure 1.14: Flowchart of the FLORIS model. The model input contains four categories: the ambient conditions, the turbine control settings, properties of the wind farm and the selection of the used submodels and the parameters therein. FLORIS maps these inputs to the fully developed three-dimensional flow field and several turbine outputs.

gree. This is related to the fact that the degradation of the turbine structure is strongly related to its temporal dynamics, which are neglected or strongly simplified in steady-state models. For the same reason, steady-state models have only sparingly been used in controllers that provide ancillary grid services (e.g., Siniscalchi-Minna et al., 2019).

The most common steady-state surrogate wind farm model used in the literature is the FLOW Redirection and Induction in Steady-state (FLORIS) model, developed by the University of Colorado Boulder, the Delft University of Technology (Gebraad et al., 2014) and the U.S. National Renewable Energy Laboratory (Annoni et al., 2018b). A flowchart of information of the FLORIS model is depicted in Figure 1.14.

Fundamentally, FLORIS provides a modular platform for wind farm models and readily combines a number of submodels from the literature. Generally, a single-wake deficit submodel (Bastankhah and Porté-Agel, 2016; Jensen, 1983) is combined with a single-wake deflection submodel (Bastankhah and Porté-Agel, 2016; Jiménez et al., 2010), a

wake summation submodel (Katic et al., 1987) and finally a turbine submodel to map the estimated flowfield to an estimated power value. Some single-wake submodels require an additional turbulence model (Crespo and Hernández, 1996).

Most commonly, the Gaussian wake deficit and wake deflection submodels inspired by Bastankhah and Porté-Agel (2016) are employed in combination with the wake summation model from Katic et al. (1987). The model from Crespo and Hernández (1996) is employed to model the turbine-induced turbulence in the wind farm. The Gaussian wake profiles are clearly seen in Figure 1.13.

1.7.2. Dynamic surrogate models

In contrast to steady-state models, quasi-static and dynamic wind farm models include at least some degree of temporal dynamics, most often being wake propagation. These models may be dynamic extensions of steady-state models (e.g., Gebraad et al., 2015; Shapiro et al., 2017b), developed from the ground up as dynamic models (e.g., Boersma et al., 2018a; Soleimanzadeh et al., 2014), or derived from higher-fidelity dynamic models (e.g., Annoni et al., 2016c).

These models simulate the second-to-second wind farm dynamics and these models therefore often go paired with a higher computational cost. Commonly, control solutions are synthesized that rely on techniques from the field of control engineering such as *Kalman filtering* and *predictive control*. Since these tools are often only proven for linear and subclasses of nonlinear models, synthesizing a reliable and stable control solution for these models is a nontrivial problem. Moreover, due to the higher sampling rate and higher computational cost than steady-state models, the computational efficiency of these controllers remains a secondary major challenge. Therefore, most of the dynamic model-based wind farm control solutions in the literature use simplified models such as ones that assume one time-invariant wind direction (e.g., Annoni et al., 2016c; Gebraad et al., 2015; Shapiro et al., 2017b) or are computationally intractable as of now (e.g., Vali et al., 2017). Objectives of these controllers include power maximization and power reference tracking for ancillary grid services.

1.8. The goals and contributions of this thesis

Reflecting on the previous sections in this chapter, it becomes clear that wind farm control encompasses many topics of research. While wind farm control for power maximization is of great interest to both the academic world and industry, adoption by the industry has been limited. The reasons for this, including the scientific gap, are presented in Section 1.8.1. The thesis at hand addresses this scientific gap to further mature the concept of wind farm control, increasing its relevance for practical implementation. The goals of this thesis are formalized in Section 1.8.2. Finally, Section 1.8.3 presents the outline of this thesis.

1.8.1. The scientific gap

Over the past decade, yaw-based wake steering in wind farms for power maximization has matured from an academic concept towards an industrial product, initially being

tested in high-fidelity simulation, then in wind tunnel experiments and finally in the field. However, while the field experiments from Fleming et al. (2017b) and Fleming et al. (2019) show situational increases in the power production, the concept is only tested for two-turbine interactions under specific inflow conditions. The question of how much wake steering increases the annual power production of a commercial wind farm thereby remains unanswered.

Furthermore, Howland et al. (2019) have performed wake steering for a full year of wind farm operation at a commercial site. While the authors show similar situational gains as what is found in the literature, the gain in the annual energy production is estimated to be negligible, thereby calling the success of the wind farm control solution into question. Paired with its situational successes, the algorithm apparently leads to losses in other situations to outcancel the measured gains in energy yield.

1.8.2. The goals of this thesis

This thesis addresses this scientific gap by moving towards the closed-loop model-based wind farm control architecture. At large, the goal of this thesis is formalized as:

Thesis objective:

Maturing wind farm control technologies for power maximization in a model-based closed-loop framework towards real-world practical applicability.

More specifically, first, the issues with the state-of-the-art wind farm control solutions for real-world implementation are highlighted. Therefore, the first research question in this dissertation is formalized as:

Research question I:

What are the limitations of the current open-loop wind farm control solutions?

Furthermore, the losses seen with the existing control solution in Howland et al. (2019) are hypothesized to be twofold, addressed by two separate research questions, respectively. Firstly, the steady-state inflow conditions and flow behavior are hypothesized to not be estimated accurately in the surrogate wind farm model, leading to excessive or unnecessary yaw misalignment angles. This eventually leads to a loss in power production. Therefore, the second research question of this thesis is formalized as:

Research question II:

Can we establish resilience to surrogate model uncertainty in modern wind farm control solutions through real-time measurement feedback and if so, how?

Consequently, this thesis introduces real-time model adaptation solutions for these existing wind farm control algorithms (leveraging steady-state surrogate wind farm models), in pursuit of reliable and accurate yaw-based wake steering solutions for power maximization.

The second hypothesized reason for the losses in Howland et al. (2019) originates from the fact that the used surrogate model is steady-state and thereby ignores the temporal dynamics. As models neglect these temporal dynamics, which turn out to play a crucial role in capturing the wind farm behavior, the model's accuracy diminishes. This leads to poor performance in wind farm control solutions that employ such models. Correspondingly, the third research question of this thesis is formalized as:

Research question III:

What benefits do dynamic surrogate models have for wind farm control and how can they be implemented?

By addressing these goals, the thesis at hand aims to mature the concept of wind farm control and stimulate its practical applicability.

1.8.3. Thesis outline

This dissertation presents four contributions to the literature to mature the concept of wind farm control for practical application. Accordingly, the outline of this thesis is as follows.

Chapter 2 presents the outcomes of a field experiment for wake steering at a commercial onshore wind farm in Italy using an open-loop model-based wind farm controller. This chapter demonstrates the strong need for further research on controller synthesis, notably on the topic of model uncertainty and how to deal with this.

Chapter 3 presents a novel theoretical analysis for real-time model adaptation using the steady-state surrogate wind farm model FLORIS in a closed-loop setting. This analysis provides fundamental insights into which measurements are necessary for the accurate reconstruction of the current ambient conditions in the wind farm when using a steady-state surrogate model.

Chapter 4 employs this theoretical analysis for the synthesis of a closed-loop wind farm control solution using FLORIS. In this controller, the surrogate wind farm model is continuously adapted to the current conditions inside the farm using measurements from the turbines. As a first of its kind in the literature, this algorithm is tested in high-fidelity under time-varying inflow conditions to stress-test the control solution.

Chapter 5 shifts its focus towards the usage of dynamic surrogate models in closed-loop model-based control solutions, since important temporal dynamics are ne-

glected in steady-state models. More specifically, this chapter presents and validates a computationally efficient model adaptation solution for a dynamic surrogate model, fusing modern tools from the field of weather forecasting and control engineering.

Each chapter contains its own introduction, disseminating the relevant literature. Moreover, each chapter contains a conclusion for that contribution. The dissertation itself will be concluded in Chapter 6, reflecting back on the contributions addressed here and outlining future research directions.

Disclaimer: The contents of Chapters 2 to 5 are largely identical to the scientific articles upon which they are based, making the chapters read somewhat as separate documents. While this makes it easier for readers to understand the topics treated in this dissertation separately, it also means that information and figures may be repeated between chapters. Moreover, literature studies in the respective chapters may be outdated, notably for Chapter 5, which was published in 2018. To ensure close resemblance between the articles published in scientific journals and the dissertation at hand, it was chosen to only apply minor changes to the work in this dissertation, improving readability and removing mistakes, rather than reiterating the work itself. This must be kept in mind when reading the dissertation.

2

TESTING OPEN-LOOP WAKE STEERING AT A COMMERCIAL ONSHORE WIND FARM

ABSTRACT

Recent field trials described in the literature demonstrate the real potential of wake steering on commercial wind farms, but also show that wake steering does not yet consistently lead to an increase in energy production for all inflow conditions. Moreover, a recent survey among experts shows that validation of the concept remains the largest barrier for adoption currently. In response, this chapter presents the results of a field experiment investigating wake steering in three-turbine arrays at an onshore wind farm in Italy. This experiment was performed as part of the European CL-Windcon project. The measurements show increases in power production of up to 35% for two-turbine interactions and up to 16% for three-turbine interactions. However, losses in power production are seen for various regions of wind directions too. In addition to the gains achieved through wake steering at downstream turbines, more interesting to note is that a significant share in gains are from the upstream turbines, showing an increased power production of the yawed turbine itself compared to baseline operation for some wind directions. Furthermore, the surrogate wind farm model, while capturing the general trends of wake interaction, lacks the details necessary to accurately represent the measurements. This chapter supports the notion that further research is necessary, notably on the topics of wind farm modeling and experiment design, before wake steering will lead to consistent energy gains in commercial wind farms.

The contents of this chapter are published as a separate research article in the *Wind Energy Science* journal, as:

B M Doekemeijer, S Kern, S K Kanev, B Salbert, J Schreiber, F Campagnolo, C L Bottasso, S Maturu, S Schuler, F Wilts, T Neumann, G Potenza, F Calabretta, F Fioretti, and J W van Wingerden. Fields experiments for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science Discussions*, pages 1–22, 2020a

2.1. Introduction

Over the last years, the concept of wake steering in wind farms has gained significant popularity in the literature (Boersma et al., 2017; Kheirabadi and Nagamune, 2019). Fundamentally, wake steering leverages the principle that intentional yaw misalignment of a wind turbine displaces its downstream wake. Thus, by choosing the right yaw misalignment, the wake formed by an upstream turbine can be directed away from a downstream turbine at the cost of a small reduction in its own power production. Consequently, this concept enables a net increase in the power production of downstream turbines and, at large, wind farms. In high-fidelity simulations, wake steering strategies are shown to increase the wind-farm-wide power production by 15% for wake-loss-heavy situations (e.g., Gebraad et al., 2016). Moreover, wind tunnel experiments indicate increases in the wind farm's power production of up to 4 – 12% for two-turbine arrays (Adaramola and Krogstad, 2011; Bartl et al., 2018; Schottler et al., 2016), up to 15 – 33% for three-turbine arrays (Campagnolo et al., 2016a,b; Park et al., 2016) and up to 17% for a five-turbine array (Bastankhah and Porté-Agel, 2019). However, these experiments neglect realistic wind variability and measurement uncertainty. A field experiment of wake steering in a scaled wind farm by Wagenaar et al. (2012) is inconclusive compared to baseline operation. In response, there has been a surge in the interest towards the development of reliable wake steering solutions that address issues of wind variability and measurement uncertainty (e.g., Chapter 4 of this dissertation and Kanev, 2020; Rott et al., 2018; Simley et al., 2020).

A small number of articles focus on the validation of wake steering at full-scale turbines and commercial wind farms. Fleming et al. (2017a) instrumented a GE 1.5MW turbine with a Light Detection And Ranging (LiDAR) system and operated the turbine at various yaw misalignments to study the wake deflection downstream. Then, Fleming et al. (2017b) demonstrated wake steering at an offshore commercial wind farm with relatively large turbine spacing of 7 to 14 times the rotor diameter ($7-14D$). These field trials involved yawing an upstream wind turbine and investigating the change in power production at the downstream turbine. When looking at two turbine pairs spaced $7D$ and $8D$ apart respectively, a gain was seen in the power production of the second turbine for most wind directions, at the cost of a much smaller loss on the upstream machine. This led to an increase in the combined power production of up to 10% for various wind directions. No significant improvements were seen for third turbine pair spaced at $14D$. However, the uncertainty bounds remain fairly large and the results also suggest that the net energy yield reduces due to wake steering for a smaller number of cases. Thereafter, Fleming et al. (2019, 2020) evaluated wake steering at a closely spaced ($3-5D$) onshore wind farm surrounded by complex terrain, again considering two-turbine interactions. Measurements show that the net energy yield can increase by up to 7% and reduce by up to the same amount for the $3D$ -spaced turbine pair, depending on the wind direction. Similarly, the change in the net energy yield for the $5D$ -spaced turbine pair is between +3% and -2.5%. It must be noted that the situations that lead to an increase in power production outnumber those that show a decrease in power production. Furthermore, Howland et al. (2019) assessed the concept of wake steering on an onshore 6-turbine

wind farm with 3.5D turbine spacing. While significant gains in power production of up to 47% for low wind speeds and up to 13% for higher wind speeds are reported for particular situations, the authors also state that the net energy gain of the wind farm over annual operation is negligible compared to baseline operation.

Considering the current literature on wake-steering field experiments, it is apparent that wake steering has real potential to increase the net energy production in wind farms, yet does not consistently lead to an increase in power production for all inflow conditions. Moreover, only Howland et al. (2019) address multiple-turbine interaction, rather than the two-turbine interactions addressed in Fleming et al. (2017b, 2019, 2020). Clearly, additional research and validation is necessary for the industry-wide adoption of wake steering control algorithms for commercial wind farms. This is in agreement with a recent survey among experts in academia and industry working on wind farm control (van Wingerden et al., 2020), which shows that the lack of validation is currently the primary barrier preventing implementation of wind farm control.

In this regard, this chapter presents the results of a field campaign for wake steering at an onshore wind farm with complex terrain in Italy, as part of the European CL-Windcon project (European Commission, 2020). The goal of this experiment is to assess the potential of the current wake steering strategies for such complicated, commercial wind farms. The contributions of this chapter are:

- As one of the few in the literature, demonstrating the potential of a state-of-the-art wind farm control algorithm for wake steering at an commercial onshore wind farm with complex terrain.
- Investigating wake interactions in non-aligned (i.e., not in a straight line) three-turbine arrays, in which yaw misalignments are applied to the first two turbines. The yaw misalignments are computed offline, based on the optimization of a simplified mathematical model of the wind farm. Wake steering for non-aligned turbine arrays has not been treated in the existing field experiments.
- The assigned yaw misalignment covers both negative and positive angles, depending on the wind direction. In the existing literature, turbines were only misaligned in one direction.
- Addressing multiple turbine types. Namely, the second turbine, WTG E5, has a different hub height and rotor diameter than the other turbines. This has not yet been assessed in the existing field experiments.

The chapter is structured as follows. Section 2.2 outlines the wind farm and the experiment. Section 2.3 shows the turbine control setpoints, calculated using state-of-the-art wind farm control solutions. Section 2.4 describes the data post-processing. Section 2.5 presents the results of the field experiment. Finally, the chapter is concluded in Section 2.6.



Figure 2.1: Positions of the wind turbines used in the wake steering campaign. Turbines WTG 26 and E5 are operated at a yaw misalignment to steer the wakes away from downstream turbines WTG E5, 10, 11, 12 and 31. WTG 25 is used for normalization. WTG E5 is a GE 1.5sle turbine and all others are GE 1.5s turbines. Imagery ©2020 Google, Imagery ©2020 CNES / Airbus, Maxar Technologies, Map data ©2020.

2.2. Methodology

This section outlines the details of the experiment. In Section 2.2.1, the wind farm layout, terrain and turbine properties are depicted. Then, Section 2.2.2 addresses the wake steering experiment itself and discusses several challenges faced compared to previous field tests. Finally, Section 2.2.3 describes what data is collected during the experiment.

2.2.1. The wind farm

The wake-steering field campaign has been executed on a subset of turbines in a commercial, onshore wind farm near Sedini on the island of Sardinia, Italy. The field experiment is part of the European CL-Windcon project. The wind farm, owned and operated by ENEL Green Power (EGP), is typically operated for commercial purposes, not for testing. EGP is a global leader in the green energy sector with a managed capacity of around 46 GW across a generation mix that includes wind, solar, geothermal and hydropower and is at the forefront of integrating innovative technologies into renewable power plants. The wind farm contains a total of 43 GE wind turbines, of which 36 turbines are of the type GE 1.5s and 7 turbines of the type GE 1.5sle. Properties of the two turbine types found in this farm are listed in Table 2.1. The relevant subset of the wind farm layout is shown in Figure 2.1. In the wake steering campaign, WTG E5 is of the type GE 1.5sle and all other turbines are of the type GE 1.5s.

The Sedini wind farm is located in a relatively flat area with an average elevation of 360 m to 400 m above sea level, surrounded by hills of 400 – 450 m above sea level. The site vegetation consists of scrub and clear areas. The predominant wind direction is from the west. The median wind speed is 4 – 6 m/s, depending on the season. The site has a median ambient turbulence intensity of 15 – 25% with a mean shear exponent

Table 2.1: General properties of the GE 1.5s and GE 1.5sle wind turbines

Variable	GE 1.5s	GE 1.5sle
Rated power (MW)	1.5	1.5
Cut-in wind speed (m/s)	4.0	3.5
Rated wind speed (m/s)	13.0	12.0
Rotor diameter (m)	70.5	77.0
Hub height (m)	65	80

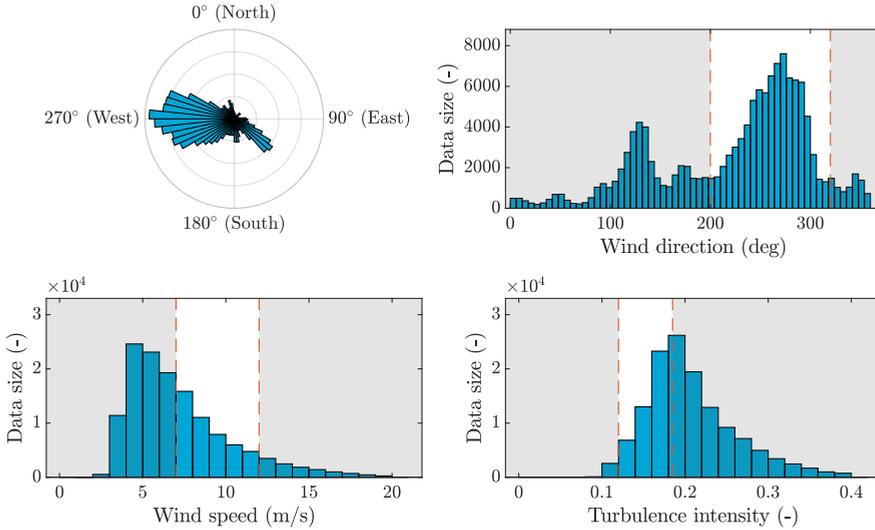


Figure 2.2: All measured data from 19 August 2019 until 3 February 2020, binned by wind direction, wind speed and turbulence intensity. Wind comes predominantly from the west, which is within the scope of the wake steering experiment. Furthermore, wind speeds are relatively low and turbulence intensities are high. The grey area covers data that is discarded in analysis of the wake steering experiments.

of 0.05 to 0.25 for day and night, respectively (S Kern et al., 2017). Figure 2.2 shows the estimated wind direction, wind speed and turbulence intensity of the data collected by the upstream turbines.

2.2.2. Experiment design

For the wake steering experiments, eight turbines are used: WTG 10, 11, 12, 24, 25, 26, 31 and E5, as shown in Figure 2.1. The situations of interest are when WTG 26 sheds a wake on WTG E5 and one or both turbines shed wakes on turbines WTG 10, 11 or 12. Additionally, for north-west wind directions, the situation where turbine WTG E5 sheds a wake on WTG 31 is of interest. For all situations, WTG 25 is used as a reference turbine and WTG 24 and WTG 25 are used to estimate the inflow ambient conditions for WTG 26 and WTG E5. While this layout lends itself well to wake steering, this field campaign faces several challenges, namely:

Table 2.2: Wind turbines of interest, scheduled according to the wind direction. To maximize the benefits of wake steering, only three turbines are considered at a time, depending on the ambient wind direction.

Wind direction	Turbines of interest
$< 235^\circ$	WTG 26, WTG E5 and WTG 10
$235^\circ - 253^\circ$	WTG 26, WTG E5 and WTG 11
$253^\circ - 276^\circ$	WTG 26, WTG E5 and WTG 12
$\geq 276^\circ$	WTG 26, WTG E5 and WTG 31

- Part of the experiment is in late summer, with higher turbulence levels and lower wind speeds compared to winter. Moreover, onshore wind farms typically experience a higher turbulence intensity than offshore farms. Higher turbulence levels generally yield lower benefits for yaw-based wake steering (Appendix A).
- There are variations in the terrain, turbine hub heights and turbine rotor diameters throughout the wind farm. Due to its high level of complexity, surrogate wind farm models address these effects to a very limited degree and lack validation with higher-fidelity and experimental data.
- The downstream turbines are closely spaced, implying that gains due to wake steering are hardly noticeable when considering the complete downstream array. For example, if the wake of WTG E5 is redirected away from WTG 10, then the combined net gain of WTG 26, E5, 10, 11, 12 and 31 would be relatively small. In addition, wake steering should be very precise, as the wake must be redirected in between WTG 10, 11, 12 and 31 to lead to a net energy increase. For example, if the wake is deflected away from WTG 11, it may be moved on top of WTG 10 or 12, thereby effectively leading to zero net gain.
- The ambient conditions are to be estimated using existing turbine sensors, rather than external measurement equipment such as a LiDAR system. This is likely to be less accurate but more realistic for the future commercialization of wake steering.

These challenges, in addition to common challenges such as irregular turbine behavior and measurement uncertainty, have led to the decision to consider only one of the downstream turbines (WTGs 10, 11, 12, 31) at a time, scheduled according to the ambient wind direction, as listed in Table 2.2. Thus, the remaining downstream turbines are ignored in the analysis. This means that the wake can be steered away from the considered turbines and onto the ignored turbines. This is exemplified in Figure 2.3, depicting what wake interactions are considered per wind direction.

2.2.3. Data acquisition

The benefit of wake steering strongly depends on the ambient conditions. Therefore, it is important to accurately characterize these inflow conditions. In this field campaign, data is acquired from a number of sources. A met mast with a height of 63.5 m is installed 0.5 km north of WTG 25. The met mast provides information about the wind speed, wind

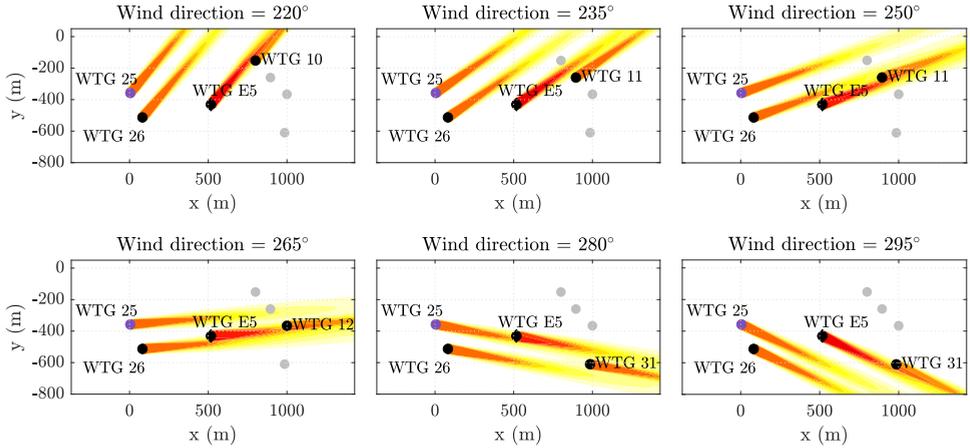


Figure 2.3: Predicted flow fields for various wind directions in baseline operation. To maximize the benefits of wake steering, only three turbines are considered at a time, depending on the ambient wind direction. The considered turbines are WTG 26, WTG E5 and one of the downstream turbines (operated without yaw misalignment). The schedule of which turbines are considered is listed in Table 2.2.

direction, vertical shear, temperature and humidity in the wind farm. However, ambient conditions vary significantly throughout the farm, not in the least due to this being an onshore wind farm. For this reason, a mobile, ground-based vertical LiDAR system of the type Leosphere WindCube v2 is installed to measure the inflow at WTG 26 for the first several months of the wake steering field campaign. This LiDAR system measures the wind speed at a 0.1 m/s accuracy and the wind direction with a 2° accuracy at 12 programmable heights up to 200 m, with a sampling rate of 1 Hz. This LiDAR system cannot communicate with the control algorithm in real time and thus was only used in postprocessing to validate the ambient wind speed estimated in front of WTG 26 using WTG 24 and WTG 25. The validation is shown in Figure 2.4, displaying a good fit.

In addition to the LiDAR system, WTG 26 and WTG E5 are instrumented with an additional, accurate nacelle anemometer. Also, WTG 12, 26 and E5 are each instrumented with an additional, accurate nacelle position sensor. Note that these sensors were only available during the first months of the field experiment, used for calibration and monitoring. The GE wind turbines provide standardized SCADA data such as the generator power, the wind speed measured by the anemometer, the wind direction measured by the wind vane and the yaw orientation measured with the yaw sensor. An algorithm internal to the GE turbine provides estimates of the 1-minute-averaged wind speed, 1-minute-averaged wind direction and 10-minute-averaged turbulence intensity.

2.3. Controller synthesis

As the research field in wind farm control is quickly evolving, an increasing amount of focus is put on closed-loop wind farm control solutions (Doekemeijer et al., 2019b). However, implementing and testing such a closed-loop wind farm control algorithm is not

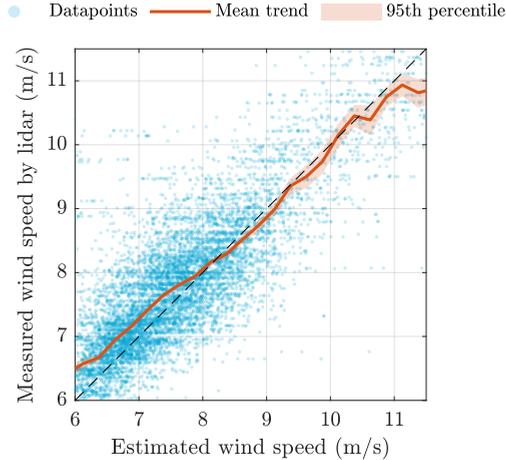


Figure 2.4: Comparison of wind speed estimates from the LiDAR and from the turbine anemometers. For the field campaign, the freestream wind speed at WTG 26 is estimated using upstream turbines WTG 25 and WTG 24. This approach is validated by comparing the estimates to measurements of the Leosphere WindCube v2 LiDAR, installed in front of WTG 26 throughout the first several months of the field campaign. The figure shows that the estimates largely match the measurements and the 95% uncertainty bounds, denoted by the shaded region, are narrow.

feasible for the designated field campaign and instead an open-loop wind farm control solution is opted for. Closed-loop solutions require additional communication infrastructure compared to open-loop solutions. Also, the actual turbine behavior becomes less predictable as the complexity of the controller increases significantly.

The controller consists of two components. Firstly, the ambient conditions are estimated, as the optimal turbine yaw setpoints vary with the inflow conditions, of which the wind direction is the most important variable. How the ambient conditions are estimated is described in Section 2.3.1. Secondly, the optimal turbine yaw setpoints for WTG 26 and WTG E5 are assigned to the turbines from a look-up table. The synthesis of this look-up table is outlined in Section 2.3.2.

2.3.1. Estimation of the ambient conditions

As outlined in Section 2.2.3, the ground-based LiDAR cannot be used in real-time for the wind farm control solution. Moreover, the met mast is located too far away to give a reliable estimate of the ambient conditions. Therefore, turbine SCADA data is used to derive an averaged freestream wind speed, wind direction and turbulence intensity for WTG 26 and WTG E5. For this purpose, the individual estimates from turbines WTG 24 and WTG 25 are averaged, which operate in freestream flow for the wind direction range considered for the wake steering experiments.

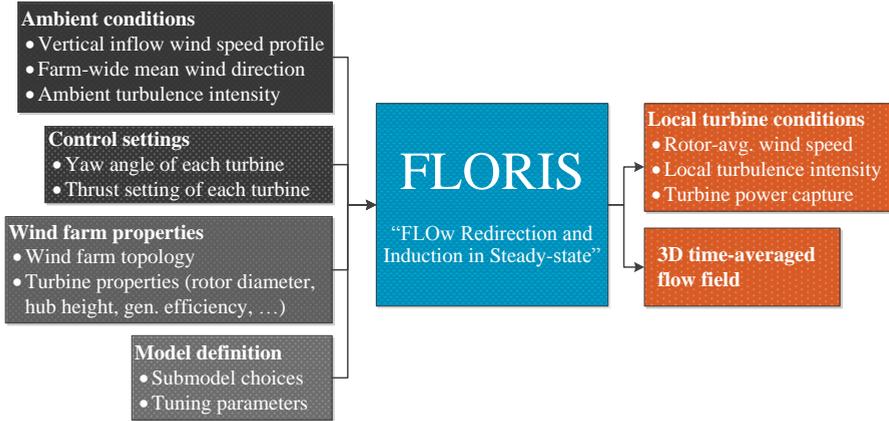


Figure 2.5: Flowchart of the FLORIS model. This model has four classes of inputs: the ambient conditions, a set of model parameters, the turbine control settings and the wind farm properties (e.g., layout). FLORIS maps these inputs in a static fashion to a set of turbine outputs being the power capture and the three-dimensional flow field.

2.3.2. Optimization of the turbine control setpoints

The turbine yaw angles are optimized using the FLOW Redirection and Induction in Steady-state (FLORIS) surrogate wind farm model, developed by CU Boulder, NREL and the Delft University of Technology (Doekemeijer and Storm, 2019; Gebraad et al., 2016). FLORIS is a surrogate wind farm model that combines several submodels from the literature, such as the single-wake model from Bastankhah and Porté-Agel (2016), the turbine-induced turbulence model by Crespo and Hernández (1996) and the wake superposition model by Katic et al. (1987). The surrogate wind farm model predicts the steady three-dimensional flow field and turbines’ operating conditions of a wind farm under a predefined inflow at a low computational cost in the order of 10 ms to 1 s. Figure 2.5 shows a flowchart of the inputs and outputs of FLORIS.

The yaw angles of WTG 26 and E5 were optimized in FLORIS for a range of wind directions (200° to 320° in steps of 2°), wind speeds (3 m/s to 13 m/s in steps of 1 m/s) and turbulence intensities (7.5%, 13.5% and 18.0%). This took approximately 10^2 CPU hours. The yaw angles are fixed between wind speeds of 5 m/s and 11 m/s in postprocessing to reduce yaw actuation at a negligible loss in the expected gains (Kanav, 2020). From wind speeds 5 m/s and 11 m/s, the angles are interpolated linearly to $\gamma = 0^\circ$ at 3 m/s and 13 m/s, respectively. This is to avoid undesirable behavior near cut-in and rated operation.

Furthermore, to reduce sensitivity of the optimized yaw setpoints to the wind direction, a Gaussian smoothing kernel was applied to the table of optimized setpoints with a standard deviation of 1.5° . The resulting look-up table for a turbulence intensity of 7.5% is shown in Figure 2.6. This figure also shows the predicted gains in power capture for the specified subset of turbines according to FLORIS in idealized conditions. It is seen

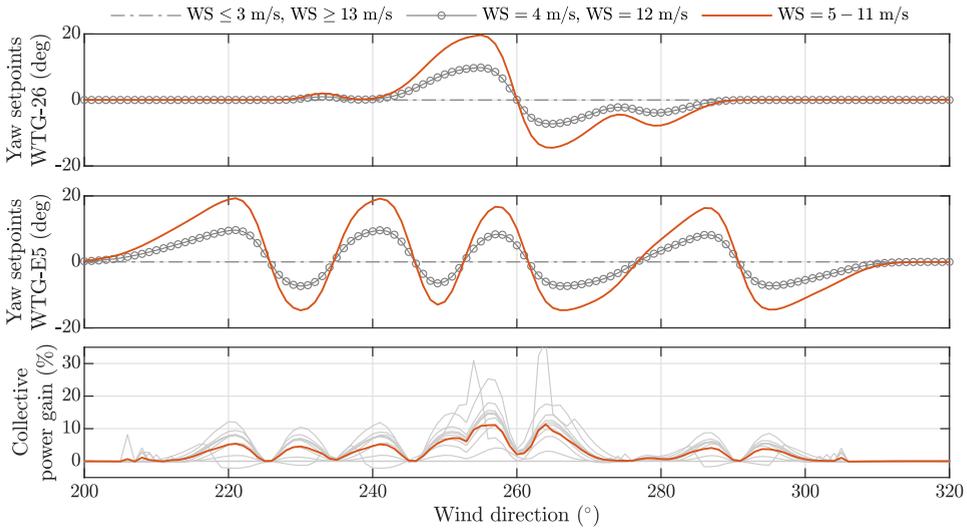


Figure 2.6: The turbine yaw setpoints for WTG 26 and WTG E5 for a freestream turbulence intensity of 7.5%. The yaw angles hold constant values for wind speeds of 5 m/s to 11 m/s. At lower respectively higher wind speeds, the setpoints are interpolated to $\gamma = 0^\circ$ at 3 m/s and 13 m/s. The collective power gain of WTG 26, WTG E5 and the downstream turbine (WTG 10, 11, 12, or 31) averaged over all wind speeds is shown as the orange line in the bottom plot. The gray lines therein represent the predicted gains for one wind speed.

that gains of 5% to 15% are expected near the wind directions 255° and 265° at a turbulence intensity of 7.5%. Furthermore, smaller gains in the order of 5% can be expected for wind directions 220° , 230° and 240° at a turbulence intensity of 7.5%. The look-up tables for higher turbulence intensities are included in Appendix A and indicate a strong decrease in expected gains for higher turbulence intensities.

FLORIS makes compromising assumptions about the wind farm terrain and wake behavior. Thus, these predictions hold a high uncertainty. As a first step to check its robustness, the optimized yaw angles from FLORIS are simulated in FarmFlow, the in-house wind farm model of TNO (Kanev et al., 2018). FarmFlow is of the same fidelity of FLORIS, but has a different set of underlying equations and therefore provides different predictions. While FarmFlow predicts lower gains, which is a common trend for FarmFlow compared to FLORIS, it also predicts little to no losses compared to baseline operation for most table entries, thereby solidifying confidence in the synthesized table of setpoints. Furthermore, after implementation in the real wind farm, the presented control module is toggled on/off every 35 minutes to allow a comparison of wake steering with baseline operation.

2.4. Data processing

Sections 2.2 and 2.3 outlined the steps taken prior to the experiment. This section now addresses how the data is processed after the experiment. One-minute-averages of SCADA

data are collected from August 19th, 2019 onward. Analysis was performed on data up until February 3rd, 2020. The data is postprocessed to eliminate any faulty or irrelevant entries as follows:

1. All data with SCADA-based wind direction estimates outside of the region of interest (200° to 320°) is discarded.
2. All data with SCADA-based ambient wind speed estimates lower than 7 m/s and higher than 12 m/s is discarded, because of high noise levels and/or the optimized yaw angle setpoints being very small in these regions (Figure 2.6).
3. All data with SCADA-based turbulence intensity estimates lower than 12% and higher than 18.0% are discarded. The upper bound is because a high turbulence intensity reduces wake effects and thereby the expected gains. Moreover, a narrow turbulence intensity range is desired with as many datapoints as possible for a fair and statistically sound analysis, explaining the lower bound. The turbulence intensity range is on the higher side due to the nature of the experiment. The specified bounds allow for a sufficient number of measurements such that a sound statistical analysis can be performed.
4. All data where the turbines of interest produce less than 200 kW of power are discarded, to reduce the relative variance in power and eliminate any situations in which turbines exhibit cut-in and cut-out behaviour.
5. Data within 5 minutes after a toggle change (baseline vs. optimized operation) is discarded.
6. Power measurements are time filtered using a (non-causal) moving average with a centered time horizon of 5 minutes.
7. The datasets are separated according to their operational mode: baseline and optimized. The datasets are then balanced such that for each wind direction and wind speed (in steps of 1 m/s), the number of measurements for baseline operation and optimized operation are equal. This reduces bias in the analysis for unbalanced bins.

With the filtered data, the energy ratio method from Fleming et al. (2019) is then used to calculate the gains due to wake steering. Important to note is that WTG 10 and WTG 11 are curtailed to a maximum of 500 kW for long periods of time during the measurement campaign. To prevent the elimination of this dataset, a part of the analysis is performed using the wind speed estimates of the local wind turbine controllers, rather than the generated power signals. Note that the analysis for WTG 10 and WTG 11 is exclusively done with measurements during curtailed operation, while the analysis for the other turbines relies on measurements during normal operation – curtailed and non-curtailed measurements are not mixed within bins.

Figure 2.7 shows the histograms of the postprocessed dataset, divided into *baseline* and *optimized* data. The relatively high turbulence intensity shown in this figure corresponds to gains in power production in the order of 2% to 6% according to FLORIS.

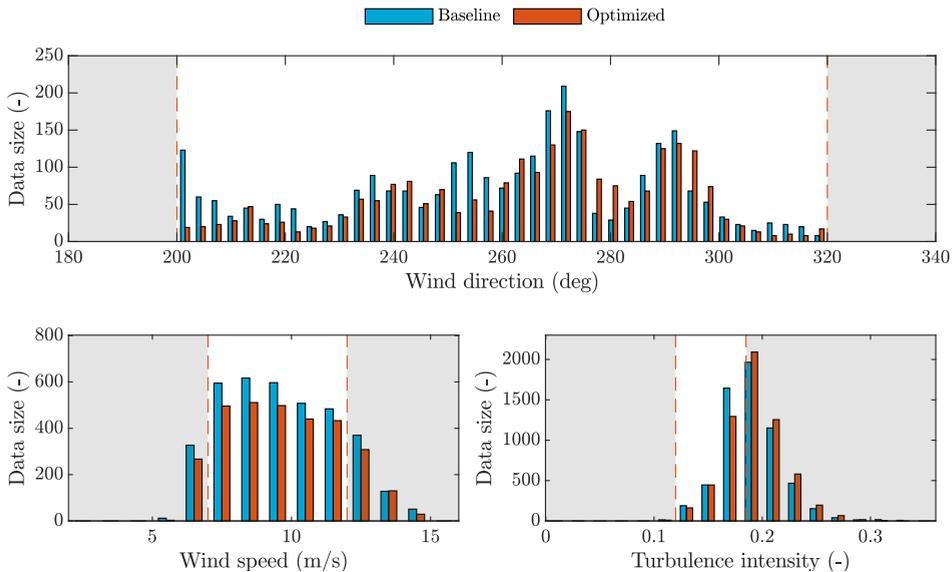


Figure 2.7: Filtered measurement data from 19 August 2019 until 3 February 2020, binned as a function of wind direction, wind speed and turbulence intensity.

2.5. Results & discussion

This section analyzes the measurement data and quantifies the change in performance due to wake steering compared to baseline operation. Note that all local wind speed estimates and power production signals are normalized with respect to the measurements from WTG 25, to reduce the sensitivity of variables to the ambient wind speed. Furthermore, 95% confidence intervals are calculated through bootstrapping (Efron and Tibshirani, 1993) for the results presented in this section.

Figure 2.8 portrays the yaw misalignment setpoints and the power production of WTG 26. The dashed lines represent the predictions from FLORIS and the solid lines represent the measurement. Since WTG 26 is not misaligned for wind directions lower than 230 degrees and higher than 290 degrees, the normalized power production should equal to 1.0, as reflected in the FLORIS predictions. Around wind directions of 255° and 265°, yaw misalignments are assigned to the turbine, expected to lead to a loss in its power production. Looking at the measurements, the yaw setpoints are successfully assigned for all wind directions. However, the predicted loss in power production due to yaw misalignment is not reflected in the measurements. Rather, it appears that positive yaw misalignment angles lead to a significant decrease of about 10% in the power production (wind directions of 240 – 250°), while negative yaw misalignment angles even lead to a slight increase in the power production compared to baseline operation (wind directions of 255 – 295°). This indicates asymmetry and a high sensitivity in the power curve for yaw misalignment, which are both not accounted for in FLORIS. These observations were confirmed with measurement data from a different GE 1.5s turbine, briefly addressed in

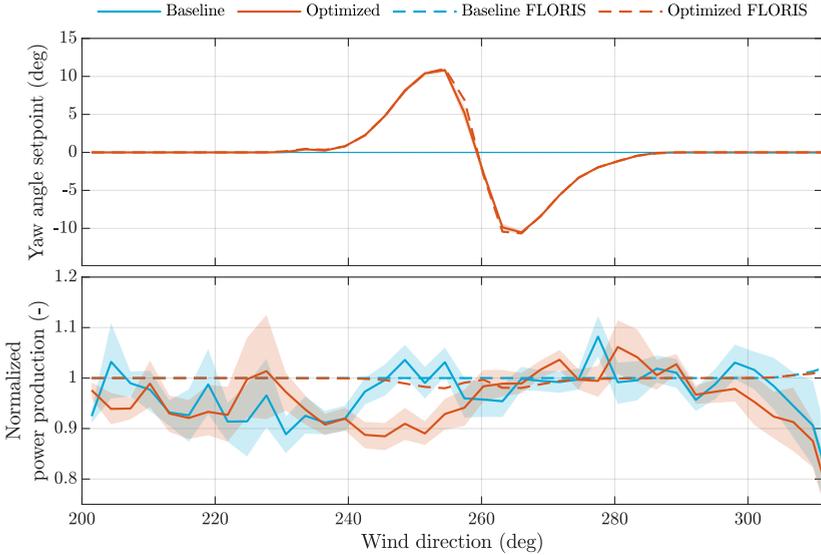


Figure 2.8: Yaw misalignments and corresponding power production for WTG 26, normalized with respect to WTG 25. The shaded areas show the 95% confidence bounds. The dashed lines represent the predictions for the measured inflow conditions by FLORIS.

Appendix B. Moreover, unknown factors lead to a systematically lower power production in the region $200 - 225^\circ$ compared to WTG 25. Also, even though both datasets operate at zero yaw misalignment in the region $295 - 320^\circ$, the *optimized* dataset shows a consistent loss compared to *baseline* operation for unidentified reasons. Hypothesized reasons for these discrepancies include terrain effects and differences in inflow conditions and turbine behavior between WTG 26 and WTG 25 to which the signals are normalized.

Figure 2.9 depicts the yaw misalignment setpoints and the power production of WTG E5. This turbine contains considerably more yaw variation between wind directions due to the close spacing and the scheduling of the considered downstream turbine (Table 2.2 and Figure 2.6). This figure shows that the yaw setpoints are applied successfully with little error. Further, note that the normalized power production for unwaked conditions is about 1.3 instead of 1.0 due to the larger rotor size and the higher tower of WTG E5. Wakes of WTG 25 and WTG 26 cause losses in power production in both *baseline* and *optimized* operation for various wind directions. These effects are both reflected in the measurements and seen in the FLORIS predictions. Notably, clear dips in the power production for both *baseline* and *optimized* operation are seen at 260° and 278° caused by wake losses. FLORIS predicts these losses, but lacks the accuracy to represent the finer trends in the measurements. Moreover, changes in the power production due to a yaw misalignment on WTG E5 appear inconsistent (e.g., large loss at 245° , no losses for 210° to 240°).

Figure 2.10 displays the cubed wind speed estimate of the downstream turbine of in-

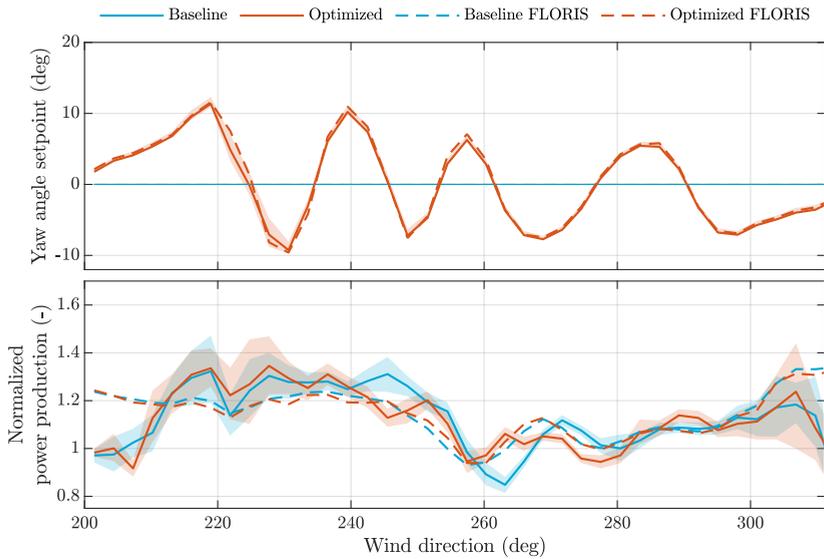


Figure 2.9: Yaw misalignments and power production for WTG E5, normalized with respect to WTG 25. The shaded areas show the 95% confidence bounds. The dashed lines represent the predictions for the measured inflow conditions by FLORIS.

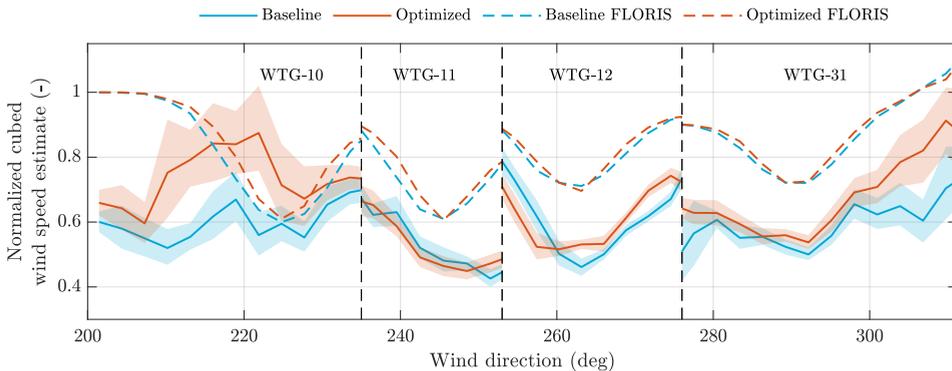


Figure 2.10: The cubed wind speed estimates of the downstream WTG of interest, serving as a surrogate for the power production under turbine derating. The results are normalized with respect to WTG 25. The shaded areas show the 95% confidence bounds. FLORIS captures the trends well, though wake losses are underestimated. Moreover, the optimized dataset appears to outperform the baseline dataset, showing a benefit due to wake steering.

terest. The reason that this variable is displayed instead of the power production is due to the fact that WTG 10 and WTG 11 are curtailed for long periods of time, rendering the power measurements unusable. FLORIS predictions show a clear trend in power production losses due to wake interactions of upstream turbines, notably at 225° , 245° , 265°

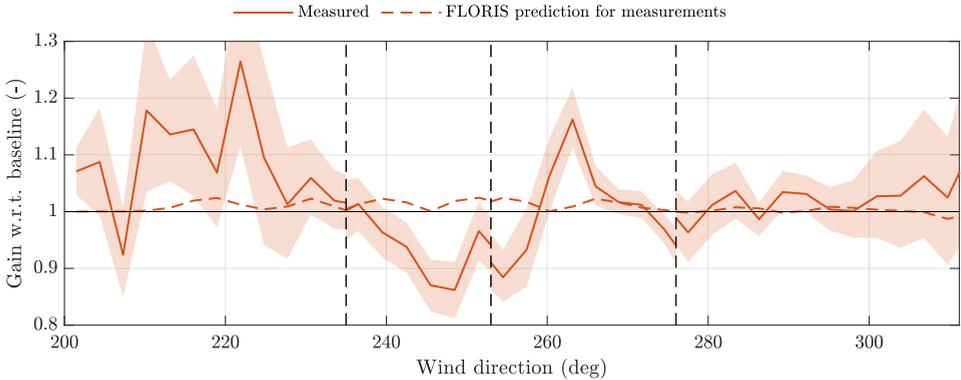


Figure 2.11: The estimated net gain of the three turbines for wake steering compared to baseline operation. The shaded area shows the 95% confidence bounds.

and 290° . Since none of the downstream turbines are yawed, FLORIS predicts that optimized operation should never lead to any losses compared to baseline operation. When looking at the measurements, these predictions are largely reflected. However, FLORIS overestimates the wake recovery and the power losses due to wake interactions are larger than predicted. This is not in the least due to the lack of an accurate terrain model. Because of the underestimated wake effects in FLORIS, wake steering should have a higher potential than predicted and the optimal yaw angles depicted in Section 2.3 may be underestimated. Moreover, the figure shows a very large increase in power production for the region $205 - 235^\circ$ between optimized and baseline operation. This is due to WTG E5 steering away its wake from WTG 10. These two turbines are positioned closest together in the wind farm and wake losses are therefore predicted to be the highest (Figure 2.1). Furthermore, gains in power production are seen in the region $260 - 320^\circ$. This somewhat agrees with where FLORIS predicts gains to be. However, the measurements also show losses near 255° . This is possibly due to the strong gradients in the yaw misalignment setpoints and thereby the sensitivity to noisy inflow conditions. Also, FLORIS predicts zero wake losses for a wind direction of 200° for both the baseline and optimized dataset, yet the measurements show a much lower wind speed. This is hypothesized to be due to topology effects and turbine interaction that was underestimated or not accounted for in FLORIS.

Finally, the change in performance for the combined three turbines is displayed in Figure 2.11. FLORIS predicts a relatively small but consistent gain across different wind directions of about 3%. This is largely due to high turbulence levels and the underestimated wake losses in FLORIS (Figure 2.10). This in turn leads to the underestimation of the benefits of wake steering. When looking at the measurements, a very large gain of up to 26% is seen at 222° . Interesting to note is that this 26% gain is the situation where WTG E5 steers its wake away from WTG 10 (Figure 2.3) and WTG 26 has no influence

on this interaction. If we only consider turbines WTG E5 and WTG 10, the combined gain in power production of turbines WTG E5 and WTG 10 is 35%. Though, it must be noted that the uncertainty bands are large for this bin. Generally, notable gains in power production are measured in the region $260 - 273^\circ$ with a gain of 16% at 263° , concerned with three-turbine interaction. Interesting to note is that all three turbines experience an increase in power production for this wind direction, be it due to a yaw misalignment or due to a steered wake. Among these three turbines, the largest gain comes from WTG E5 with a 29% increase in power by itself. Furthermore, Figure 2.11 also shows notable losses, especially in the region near 250° , due to large losses at WTG 26 originating from yaw misalignment and no gains downstream. Losses are also seen near the transition regions (black dashed vertical lines), possibly due to strong gradients in the yaw angles at these wind directions.

In addition to the mismatch between FLORIS and the actual yaw-power curve of WTG 26 and WTG E5, the lack of terrain effects in FLORIS are expected to have a significant impact on the results. This may be one of the key reasons for the overestimation of wake recovery in the FLORIS model, which in turn leads to an underestimation of the benefits of wake steering. Moreover, unmodeled effects such as secondary steering (Martínez-Tossas et al., 2019) may be a source of error. These unmodeled effects can have a positive effect on the success of wake steering. This leads to an underestimation of the potential benefits of wake steering and consequently to suboptimal yaw misalignment setpoints. Historical operational data may also be used to reduce the model-plant mismatch (Schreiber et al., 2020).

2.6. Conclusions

This chapter presented a field experiment for wake steering at a commercial onshore wind farm in Italy. Three-turbine interaction was considered, with the first two turbines operating under yaw misalignments to maximize the collective power production. The yaw setpoints were calculated according to an open-loop steady-state and model-based wind farm control solution. The field experiment shows significant gains, especially for two-turbine interaction, with an increase in combined power production of up to 35% for one particular two-turbine situation. Moreover, gains in power production for the three-turbine array up to 16% were measured for particular wind directions. However, the measurements also show notable losses for a region of wind directions, largely due to losses at the yaw-misaligned upstream turbines and due to insufficient or incorrect wake steering downstream.

Several important observations were made from the measurement data. Measurements shows that upstream turbines may benefit from nonzero yaw misalignment, already leading to an effective increase in power production at these turbine without considering the phenomenon of wake steering downstream. Such effects have a large influence on the results presented in this chapter. Moreover, the potential of wake steering was confirmed for a large range of conditions. These two factors effectively suggest that the power production in wind farms could be increased for “free”, thus allowing wake steering without losing or even increasing the energy yield upstream. Also, while the sur-

rogate wind farm model leveraged in this chapter is able to predict the dominant trends of wake interaction, large discrepancies are seen between its predictions and the field measurements. Notably, FLORIS assumes a symmetrical yaw-power curve of WTG 26 and WTG E5, assuming peak power production at zero yaw misalignment. In addition, FLORIS lacks important terrain effects and appears to overestimate wake recovery. Consequently, FLORIS underestimates the benefits of wake steering and the assigned yaw angles in this experiment are suboptimal.

At large, important lessons learned from this experiment are:

- An accurate characterization of the physical wind turbines in the surrogate model is essential. This chapter demonstrated the strong need for an accurate yaw-power curve of each turbine to maximize the benefits of wake steering and operation under yaw misalignment.
- To clearly distinguish the benefits of wake steering from baseline operation, a reliable baseline controller must be established and implemented. This may require more accurate wind direction and yaw sensors that ensure that upstream turbines accurately track the wind direction and maximize their power production.
- In this experiment, which turbine was considered to be the “downstream turbine of interest” was decided according to the wind direction to maximize the potential benefits of wake steering. Unfortunately, this is expected to be the reason for poor performance near the transition regions. Such scheduling requires more research before implementation and rather should be avoided whenever possible.
- The surrogate wind farm model is hypothesized to lack, i.a., essential temporal dynamics and complex terrain effects, leading to suboptimal yaw setpoints and controller performance. Moreover, the turbulence model in FLORIS should be improved and ideally calibrated to field data before usage.
- Field campaigns should run for at least one year to minimize the impact of measurement uncertainty. Moreover, experiments ran throughout the year will provide a realistic idea of the efficacy of the tested concept and its impact on the annual energy production.

Finally, loads are neglected in this chapter, but play a vital role in adoption of the concept (e.g., Damiani et al., 2018). In conclusion, this chapter supports the notion that further research is necessary, notably on the topic of wind farm modeling, before wake steering will lead to consistent energy gains in commercial wind farms.

3

WHAT CAN WE ESTIMATE WITH THE MEASUREMENTS AVAILABLE?

ABSTRACT

Wind farm control (WFC) algorithms rely on an estimate of the ambient wind speed, direction and turbulence intensity in the determination of the control setpoints. However, the measurements available in a commercial wind farm do not always carry sufficient information to estimate these quantities. In this chapter, a novel measure (“observability”) is introduced that quantifies how well the ambient conditions can be estimated with the measurements at hand through a model inversion approach. The usefulness of this measure is shown through several case studies. While the turbine power signals and the inter-turbine wake interactions provide information on the wind direction, the case studies presented show that there is a strong need for wind direction measurements for WFC to sufficiently cover observability for any ambient condition. Further, generally, more wake interaction leads to a higher observability. Also, the mathematical framework presented in this chapter supports the straight-forward notion that turbine power measurements provide no additional information compared to local wind speed measurements, implying that power measurements are superfluous. Irregular farm layouts result in a higher observability due to the increase in unique wake interaction. The findings in this chapter may be used in WFC to predict which ambient quantities can (theoretically) be estimated. The authors envision that this will assist in the estimation of the ambient conditions in WFC algorithms and can lead to an improvement in the performance of WFC algorithms over the complete envelope of wind farm operation.

The contents of this chapter have been published as a separate research article in the peer-reviewed *Wind Energy* journal, as:

B M Doekemeijer and J W van Wingerden. Observability of the ambient conditions in model-based estimation for wind farm control: A focus on static models. *Wind Energy*, 2020. URL <https://doi.org/10.1002/we.2495>. Early view

3.1. Introduction

The European Wind Energy Association (EWEA) predicts the amount of installed wind energy to increase from 106 GW in 2012 to 735 GW in 2050, which at that point should provide for about 50% of the European Unions electricity demand (European Wind Energy Association (EWEA), 2019). The success of wind energy largely relies on its financial competitiveness with other renewable and non-renewable sources. Control plays an invaluable role in this matter. In the past, the focus of control research has been on wind turbine control. Recently, the interest has largely shifted towards wind farm control (WFC), in which multiple turbines inside a wind farm are coordinated together to improve their combined energy yield (Boersma et al., 2017). WFC addresses the issue of wakes, which are slower and more turbulent pockets of air that form behind a wind turbine as energy is extracted. Wake formation has led up to an estimated 23% loss in the annual energy yield of the closely-spaced Lillgrund offshore wind farm at the coast of Sweden compared to an idealized situation without wake formation (Barthelmie et al., 2010). The underlying concept of WFC is to influence the wake such that it has a smaller impact on downstream turbines. A popular approach in the literature is yaw-based wake steering, in which the wake position is shifted laterally by purposely operating an upstream turbine at a yaw misalignment. Recent studies have shown the potential of yaw-based wake steering for wind farm power maximization in high-fidelity simulation (Gebraad et al., 2016), wind-tunnel experiments (Campagnolo et al., 2016a) and field tests under dynamic inflow conditions (Fleming et al., 2017a, 2019; Simley et al., 2020). These publications suggest an increase in the annual energy yield in the order of one percent and situational increases of up to twenty percent for certain wind farms under particular inflow conditions that cause large wake losses.

The wake losses and therefore the amount of yaw misalignment that maximizes the energy yield is highly dependent on the wind direction, wind speed and turbulence intensity of the incoming wind field (Barthelmie et al., 2010). As these atmospheric conditions constantly change, so do the optimal yaw angles. Typically, a simplified (“surrogate”) model of the flow and turbine dynamics is leveraged to calculate the optimal yaw angles (Doekemeijer and Storm, 2019). However, due to the complicated flow behavior at a range of temporal and spatial scales, no surrogate wind farm model exists that is accurate for all the different atmospheric conditions a wind farm may encounter. For this reason, closed-loop wind farm control solutions are becoming increasingly popular in the literature (Boersma et al., 2017). The underlying idea of this closed-loop control framework is that the surrogate wind farm model is continuously adapted such that it accurately and consistently predicts the wind farm behavior.

The closed-loop WFC framework is shown in Figure 3.1. This framework consists of three components, namely 1) a surrogate wind farm model, 2) a model adaptation algorithm and 3) a control setpoint optimization algorithm. Surrogate wind farm models can typically be separated into static and dynamic models. These model types attempt to predict the minute-averaged and the second-to-second flow and turbine behavior, respectively. The purpose of the model adaptation algorithm is to modify parameters inside the surrogate model such that it can accurately predict the wake interactions inside

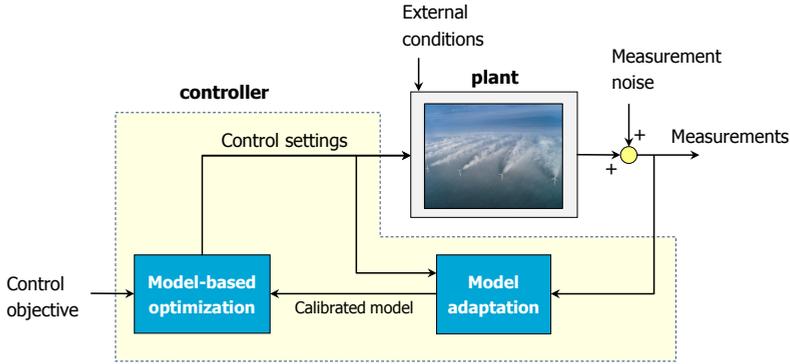


Figure 3.1: The closed-loop framework for model-based wind farm control. In this framework, measurements from each turbine in the wind farm (e.g., turbine power signals, wind vane measurements) are used to adapt a simplified model of the wind farm to better represent the current wind farm dynamics. Typically, the freestream wind speed and wind direction are among the estimated quantities. This adapted surrogate model is then used to optimize the control settings of each turbine to increase the power capture of the wind farm. Finally, these setpoints are transmitted to the real-world turbines and the cycle repeats itself.

the wind farm, which includes the freestream wind speed and wind direction. Finally, an optimization algorithm is necessary to determine an optimal control policy such that a particular wind farm objective is achieved, e.g., maximization of the wind farm power production. The focus in this chapter is on the model adaptation algorithm; the interested reader is referred to the survey by Boersma et al. (2017) for more information on surrogate wind farm models and optimization algorithms.

The body of literature on real-time model adaptation for WFC is scarce. Most WFC literature has focused on setpoint optimization and model development (Boersma et al., 2017). This goes paired with the fact that most WFC algorithms in the literature have been tested under quasi-steady ambient conditions, meaning that the mean wind speed, wind direction and turbulence intensity were time invariant. This holds for both numerical simulations (e.g., Gebraad et al., 2016) and real-world scaled experiments (e.g., Campagnolo et al., 2016a,b). This limits the applicability of such algorithms, as the experiments do not sufficiently represent the real-world fluctuations in the atmosphere.

A handful of articles in the literature is concerned with the estimation of atmospheric conditions and model adaptation for WFC. Annoni et al. (2018a) proposed a model-free algorithm to estimate the wind direction inside a wind farm using the wind vane measurements of different turbines and obtaining a consensus on the most probable value. Doekemeijer et al. (2019b) proposed a method to estimate the freestream conditions by a model inversion approach using the time-averaged turbine power measurements and a static surrogate wind farm model assuming the wind direction is known, which is comparable to the idea coined by Gebraad et al. (2016). Furthermore, Gebraad et al. (2015) synthesized a Kalman filter for their dynamic surrogate model, which uses the turbine power measurements to estimate the flow field inside the wind farm. The adapted surrogate model was able to accurately predict the wind farm dynamics, though the wind

direction was constant and assumed to be known. Similarly, Doekemeijer et al. (2018a), also presented as Chapter 5 of this dissertation, uses a dynamic surrogate wind farm model with an Ensemble Kalman filter to estimate the flow field and turbulence intensity using turbine power measurements. High-fidelity simulations show that the algorithm was able to successfully reconstruct the dynamic wind field for a 2-turbine and a 9-turbine wind farm. However, also in this work, the wind direction is assumed known. Further, Shapiro et al. (2017b) presents an Ensemble Kalman filter for wind field estimation using their simplified wind farm model, assuming a constant and known wind direction. Also, Iungo et al. (2015) used a dynamic mode decomposition method to obtain a reduced-order model of the wind farm, after which they applied a Kalman filter for state/wind field estimation. This method again assumes a known wind direction and additionally assumes the complete flow field to be measured. Besides the estimation of the ambient conditions, Bottasso and Schreiber (2018) attempt to estimate several model tuning parameters to improve the accuracy of the surrogate wind farm model.

All aforementioned work, apart from that of Annoni et al. (2018a), is tested under quasi-steady ambient conditions in simulation, thereby significantly limiting their applicability. In essence, these methods combine a set of measurements with some sort of surrogate or consensus model that relate one measurement to another. These methods are fundamentally limited due to the fact that only a finite amount of information is measured. One can easily think of situations in which the ambient conditions cannot be derived from the available measurements. Such a situation would be considered “unobservable” or “unestimable”. Thus, before one may attempt to estimate the ambient conditions, one should consider whether the situation is observable in the first place. However, to the best of the authors’ knowledge, there is no literature on the observability for ambient condition estimation. This chapter aims to fill this scientific gap and the contributions of this chapter are:

- Proposing a formal definition for a mathematical measure (henceforth referred to as “observability”) that quantifies how well the ambient conditions (i.e., wind direction, wind speed, turbulence intensity) can be reconstructed from the measurements available in the wind farm.
- Comparing the effect of different wind farm topologies and sensor configurations on the observability for a large range of ambient conditions that a wind farm may encounter during operation.
- Performing theoretical case studies with wind farms with DTU 10MW wind turbines.

This chapter is organized as follows. The surrogate wind farm model used in this chapter is presented in Section 3.2. The issue of estimation and a novel quantitative measure of observability is presented in Section 3.3. Simulation results are shown in Section 3.4 and the chapter is concluded in Section 3.5.

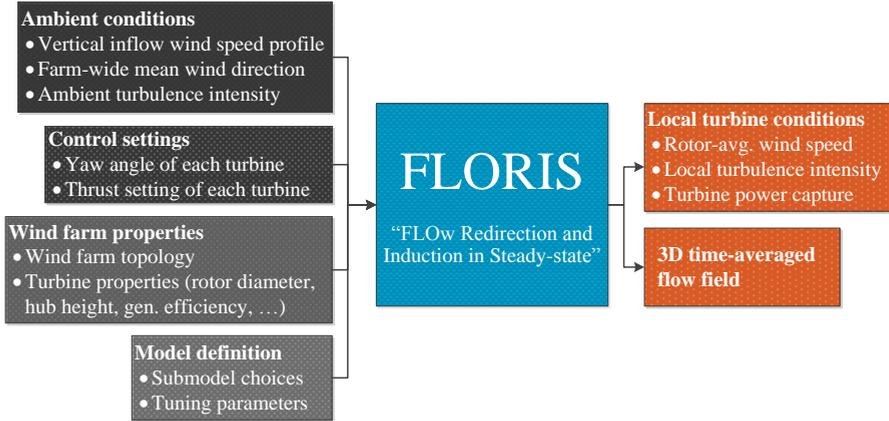


Figure 3.2: The flow of information for the surrogate model “FLORIS”. The left four blocks represent the various model inputs and the right two blocks represent the model outputs. Typically, the control settings and wind farm properties are known and are time invariant. However, the ambient conditions are time variant and the tuning parameters that provide the best results are uncertain.

3.2. Surrogate model: FLORIS

The surrogate wind farm model used in this chapter is referred to as the “FLOW Redirection and Induction in Steady-state” (FLORIS) model (Doekemeijer and Storm, 2019). This model predicts the time-averaged power capture of each turbine and the time-averaged three-dimensional flow field for a wind farm under a specified set of inflow conditions. The simulated outputs from FLORIS respond to flow and turbine quantities averaged over a span of approximately 5 minutes. A schematic overview of the types of inputs and outputs to the FLORIS model is shown in Figure 3.2. Fundamentally, FLORIS combines several submodels from the literature. The main components of FLORIS are described in the remainder of this section.

Firstly, FLORIS includes the single-turbine wake model from Bastankhah and Porté-Agel (2016), which predicts the time-averaged three-dimensional wind field behind a turbine. Secondly, the turbine-induced turbulence is calculated using an empirical function proposed by Crespo and Hernández (1996). Thirdly, the wind field under multiple overlapping wakes is calculated through a sum-of-squared-deficits law as proposed by Katic et al. (1987). Fourthly, the power production of each turbine is calculated using the rotor-effective wind speed and the nondimensional power coefficient C_P , as

$$P_i = \frac{1}{2} \rho A_D U_i^3 C_P(U_i, \gamma_i), \quad (3.1)$$

where ρ is the air density, A_D is the rotor swept area, U_i is the spatially averaged inflow wind speed at turbine i and γ_i is the yaw angle of the turbine relative to the incoming wind. The nondimensional power coefficient, C_P , can be derived using actuator disk theory for aligned inflow ($\gamma_i = 0$). Alternatively, C_P can be calculated using an aero-elastic turbine simulation model for various wind speeds (and yaw misalignment angles)

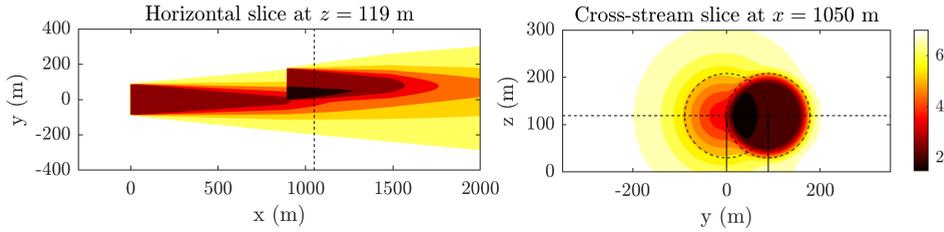


Figure 3.3: Simulation results for a wind farm with two DTU 10MW wind turbines using FLORIS. The figure shows a horizontal flowfield and a cross-stream slice of the flow-field. The Gaussian nature of the wakes is clearly seen.

such as OpenFAST (National Renewable Energy Laboratory (NREL), 2019) or Bladed. A common expression modeling the effect of a yaw misalignment on the turbine power production is (Gebraad et al., 2016)

$$C_P(U_i, \gamma_i) = C_P(U_i, 0) \cdot \cos^\kappa(\gamma_i), \quad (3.2)$$

where κ has a typical value of 1.4 – 2.0, depending on the wind turbine.

The results of an arbitrary wind farm simulation with two 10MW turbines (Bak et al., 2012) is shown in Figure 3.3. The computational cost for a single FLORIS run is 10 ms to 1 s, depending on the number of turbines in the wind farm. FLORIS has shown a good match with results from high-fidelity simulations (Doekemeijer et al., 2019b), wind tunnel experiments (Schreiber et al., 2017) and field tests (Annoni et al., 2018b; Fleming et al., 2017a). Furthermore, the variant presented in this chapter has fewer tuning parameters than a comparable model proposed in Gebraad et al. (2016). For a more detailed, mathematical description of the model, the reader is referred to its related literature. Note that the results that will be presented in this chapter are not limited to FLORIS and can straight-forwardly be reproduced with other static surrogate wind farm models.

3.3. Methodology: introducing a measure of observability

The model adaptation solution of a WFC algorithm is not guaranteed to result in satisfactory performance. There has to be sufficient information in the wind farm measurements to correctly determine the ambient conditions. Hence, an observability analysis provides useful insight before the implementation of such a control algorithm. The traditional definition of observability refers to dynamical systems; a system is observable if the initial conditions and the timeseries of the system states can be reconstructed from a timeseries of the system output signals. As FLORIS is a static model, such a notion does not apply. Therefore, a *static observability* notion is defined as being true for a situation when the initial conditions can be reconstructed from the system output signals. In this section, a mathematical definition of this “observability” is introduced for the control framework presented in Figure 3.1.

3.3.1. Cost function in estimation

Generally, a simplistic, heuristic approach is used to determine the prevailing ambient conditions inside the wind farm. However, the reliability of such methods vary, the literature on them is scarce and these methods are limited in their accuracy. Rather, in this chapter, a surrogate wind farm model is leveraged in a sensor fusion approach for the estimation of the ambient conditions.

In this chapter, the freestream wind speed, wind direction and turbulence intensity are estimated using the readily available measurements of each turbine. For example, consider a cost function that minimizes the error with the time-averaged power measurements of each turbine, as

$$J_1(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) = \frac{1}{N_T} \sum_{i=1}^{N_T} \left(P_i^{\text{measured}} - \hat{P}_i^{\text{FLORIS}}(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) \right)^2, \quad (3.3)$$

with N_T the number of turbines and $\hat{\phi}$, \hat{U}_∞ and \hat{I}_∞ being the freestream wind direction, wind speed and turbulence intensity as evaluated in FLORIS, respectively.¹ Using this cost function for model adaptation, the idea is that values for $\hat{\phi}$, \hat{U}_∞ and \hat{I}_∞ are found such that the error between the measured turbine power signals and what is predicted by FLORIS for these conditions is minimized. The cost function shown in Equation 3.3 was used for model adaptation in Doekemeijer et al. (2019b) assuming ϕ was known a priori, which allowed the successful estimation of U_∞ and I_∞ . However, only using the turbine power measurements may lead to situations in which the true ambient conditions cannot be reconstructed accurately. For example, consider the case in which all turbines inside the wind farm are operating in above-rated conditions. All turbines are then generating their rated power and one cannot distinguish different above-rated wind speeds from one another. To resolve this issue, one can include the wind speed estimates from a local turbine wind speed estimator (Ortega et al., 2013; Soltani et al., 2013) in the cost function. This term is denoted by J_2 , given as

$$J_2(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) = \frac{1}{N_T} \sum_{i=1}^{N_T} \left(U_i^{\text{measured}} - \hat{U}_i^{\text{FLORIS}}(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) \right)^2, \quad (3.4)$$

where U_i^{measured} is the measurement of the local wind speed estimator of turbine i and $\hat{U}_i^{\text{FLORIS}}$ is what FLORIS predicts the local wind speed to be at turbine i for the hypothesized wind conditions $\hat{\phi}$, \hat{U}_∞ and \hat{I}_∞ . Note that the inflow wind speed at a turbine in FLORIS, denoted by \hat{U}_i , is the freestream-equivalent wind speed at that turbine under zero yaw misalignment. Thus, the effects of a yaw misalignment of turbine i are not accounted for in this signal. However, in practice, a typical local turbine wind speed estimator provides a freestream-equivalent wind speed using the turbine power signal

¹Note that the power measurements of different turbines can be weighted differently according to the amount of uncertainty in this measurement, as done in previous work (Doekemeijer et al., 2019b). However, for the observability analysis at hand, this additional level of complexity does not sufficiently add to the theoretical foundation presented in this chapter.

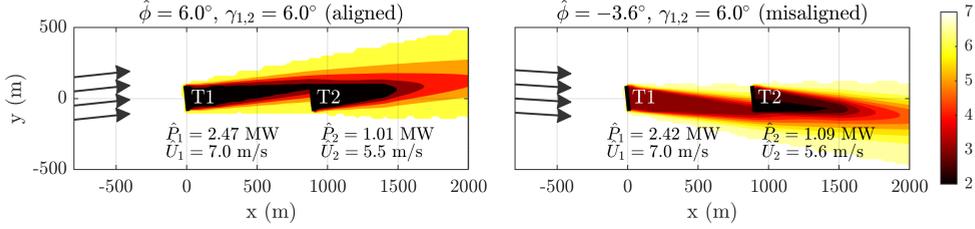


Figure 3.4: The issue of symmetry exemplified on a two-turbine wind farm for the estimation of ϕ under a fixed (in inertial frame) yaw angle γ . The definitions are that $\phi = 0^\circ$ when the air moves from west to east (left to right) and is counter-clockwise positive. The colorbar depicts wind speed in m/s^{-1} . In this plot, it is seen that $\hat{\phi} = 6^\circ$ and $\hat{\phi} = -3.6^\circ$ yield almost identical turbine power signals and local wind speeds, thus making them indistinguishable in the cost function of Equation 3.6. This leads to an unobservable situation.

under the assumption of zero yaw misalignment. To account for the situation in which a turbine is misaligned with the flow, one can model $\hat{U}_i^{\text{FLORIS}}$ as

$$\hat{U}_i^{\text{FLORIS}}(\gamma_i) = \hat{U}_i^{\text{FLORIS,unyawed}} \cdot \sqrt[3]{\cos^\kappa(\gamma_i)}, \quad (3.5)$$

in order to match the signal definition from the local wind speed estimator, U_i^{measured} . Finally, one can combine J_1 and J_2 into cost function $J_{1,2}$, as

$$J_{1,2} = \lambda_P J_1 + \lambda_U J_2, \quad (3.6)$$

where λ_P and λ_U are weighing factors. Using the cost function defined in Equation 3.6, difficult situations may arise when trying to estimate ϕ , U_∞ and I_∞ . For example, if there is no wake interaction, one cannot estimate the freestream turbulence intensity, as the effects of I_∞ have no correlation with (i.e., impact on) the measured signals. Moreover, issues may arise concerning the estimation of ϕ , as demonstrated in Figure 3.4. In this situation, $\hat{\phi} = 6.0^\circ$ and $\hat{\phi} = -3.6^\circ$ yield almost identical values for the wind speed ($\hat{U}_1^{\text{FLORIS}}$ and $\hat{U}_2^{\text{FLORIS}}$) and power production ($\hat{P}_1^{\text{FLORIS}}$ and $\hat{P}_2^{\text{FLORIS}}$), thereby making it impossible to distinguish these two situations using the measurements available.

To address the latter issue, local wind direction estimates of each turbine are included in the cost function, e.g., using the filtered wind vane measurements (Kragh and Fleming, 2012). This term is modeled as J_3 , given by

$$J_3(\phi) = \frac{1}{N_T} \sum_{i=1}^{N_T} (\phi_i^{\text{measured}} - \hat{\phi})^2, \quad (3.7)$$

where ϕ_i^{measured} is the filtered wind vane measurement of turbine i and $\hat{\phi}$ is the hypoth-

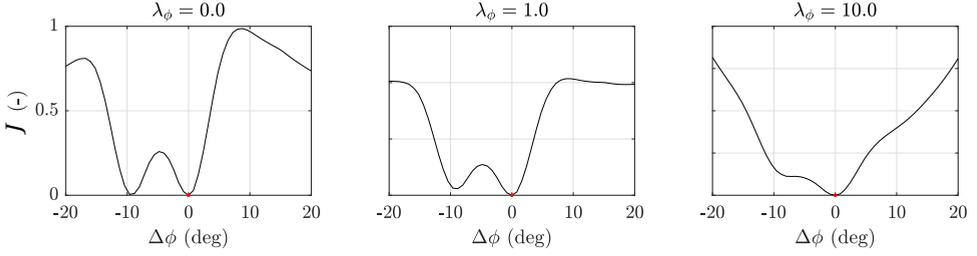


Figure 3.5: The issue of exclusively using power measurements in the cost function J (Equation 3.8), exemplified on the two-turbine case of Figure 3.4. In all subplots, $\lambda_P = 10^{-12}$ and $\lambda_U = 0$ (In this example case, $\lambda_U = 0$ as it carries the same information as the power signals do. This statement will be demonstrated in Section 3.4.1.1). In the left figure, $\lambda_\phi = 0$ and thus exclusively power measurements are used. This leads to a critical point at $\Delta\phi = -9.6^\circ$ which has negligible cost and thus this point cannot be distinguished from the actual point $\Delta\phi = 0^\circ$, with $\hat{\phi} = \phi + \Delta\phi$, leading to unobservability. This refers back to the situation shown in Figure 3.4. By including wind vane measurements ($\lambda_\phi > 0$), the cost function is better conditioned to uniquely estimate ϕ . Note that λ_ϕ should be chosen in accordance with the vane's measurement reliability.

sized wind direction in FLORIS. The complete cost function J is now defined as

$$\begin{aligned}
 J(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) = & \underbrace{\frac{1}{N_T} \sum_{i=1}^{N_T} \lambda_P \left(P_i^{\text{measured}} - \hat{P}_i^{\text{FLORIS}}(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) \right)^2}_{\text{Contribution of local power measurements}} + \\
 & \underbrace{\frac{1}{N_T} \sum_{i=1}^{N_T} \lambda_U \left(U_i^{\text{measured}} - \hat{U}_i^{\text{FLORIS}}(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty) \right)^2}_{\text{Contribution of local wind speed estimates}} + \\
 & \underbrace{\frac{1}{N_T} \sum_{i=1}^{N_T} \lambda_\phi \left(\phi_i^{\text{measured}} - \hat{\phi} \right)^2}_{\text{Contribution of local wind direction estimates}}, \quad (3.8)
 \end{aligned}$$

with λ_ϕ a weighing factor for the local wind direction estimates. This weighing factor is to be chosen according to the relative measurement noise and bias in the wind vane measurements and could vary per turbine. The to-be-estimated quantities are ϕ , U_∞ and I_∞ . Each of the three components includes a squared term to quadratically penalize mismatches between the surrogate model and sensor measurements. The situation of Figure 3.4 becomes increasingly better conditioned as the contribution of the wind vane measurements increases, as visualized in Figure 3.5.

Thus, it is clear that local wind speed measurements (to deal with above-rated wind speeds), wind direction measurements (to deal with situations as exemplified in Figure 3.4) and wake interaction (to enable correlation between I_∞ and the measurements) are required to promote observability of the freestream conditions over the full range of operation. When multiple minima exist at a notable distance from the true solution (in the example case of Figures 3.4 and 3.5 this would be $\|\Delta\phi\| \gg 0$, with $\hat{\phi} = \phi + \Delta\phi$), the

ambient conditions cannot be reliably estimated and the situation becomes “unobservable”.²

However, while it is clear that particular situations are unobservable, a quantitative measure is still required to determine the degree of (un)observability. For example, is the situation in Figure 3.5 with $\lambda_\phi = 1.0$ “observable enough” to uniquely determine the ambient conditions? To answer such questions, a quantitative measure of observability for static models is introduced in the next section.

3.3.2. A quantitative measure for observability

With the cost function defined, a quantitative measure on the degree of observability of a particular situation is defined. With “situation”, we imply a particular wind farm layout, the true ambient conditions and a specific choice of the regularization terms λ_P , λ_U and λ_ϕ . The main contribution of this chapter is the introduction of such a mathematical notion for observability. The observability of a particular situation \mathcal{O} is defined as

$$\mathcal{O} = \min(\mathcal{M}), \quad (3.9)$$

$$\text{where } \mathcal{M} = \begin{cases} \infty & \text{if } \|\Delta\phi\| < b_\phi, \|\Delta U_\infty\| < b_U, \|\Delta I_\infty\| < b_I, \\ \mathcal{F} & \text{otherwise,} \end{cases} \quad (3.10)$$

$$\text{and } \mathcal{F} = \frac{J(\hat{\phi}, \hat{U}_\infty, \hat{I}_\infty)}{k_\phi(\Delta\phi)^2 + k_U(\Delta U_\infty)^2 + k_I(\Delta I_\infty)^2}, \quad (3.11)$$

with J as defined in Equation 3.8, k_ϕ , k_U and k_I denoting normalization terms and b_ϕ , b_U and b_I being thresholds. Further, $\Delta\phi = \phi - \hat{\phi}$, $\Delta U_\infty = U_\infty - \hat{U}_\infty$ and $\Delta I_\infty = I_\infty - \hat{I}_\infty$ denote the difference between the true and hypothesized ambient conditions, respectively. In the remainder of this section, the working principle will be explained.

The function \mathcal{M} is defined such that critical points (low cost J , far away from the true solution) have a low value (less observable – hard to tell apart from the true solution), while situations in which the cost J is high yields a high value (more observable – easier to distinguish from the true solution). Furthermore, the threshold terms are present to ensure that any value estimated close enough to the true optimum does not “endanger” the observability. A more elaborate discussion on these thresholds can be found in Appendix C.

Figure 3.6 demonstrates how the observability \mathcal{O} is calculated for the example situation discussed in Section 3.3.1, with $\lambda_\phi = \lambda_U = 0$. Note that this is not necessarily a realistic scenario, but rather is discussed to provide insight into the method. The function \mathcal{M} is derived from the cost function J following Equations 3.10 and 3.11. The cost function has two minima: one at $\Delta\phi = -9.6^\circ$ and one at $\Delta\phi = 0^\circ$, indicating that there are two hypothetical wind directions that produce near-identical turbine power signals. This leads to a low observability.

²Note that observability has a different notion in the field of control engineering for dynamical systems. In this thesis, an equivalent definition is defined for the static problem outlined in this section.

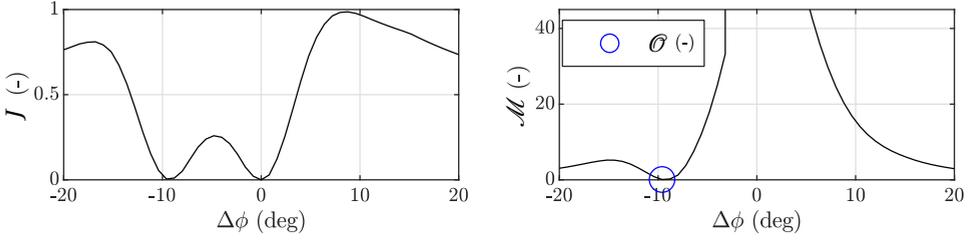


Figure 3.6: A visualization of how the degree of observability Θ is calculated. This is a continuation of the example shown in Figures 3.4 and 3.5, with $\lambda_\phi = \lambda_U = 0$. Firstly, the cost function J (left plot) is converted to a measure \mathcal{M} (right plot) which penalizes a low cost far away from the true solution (the true solution being $\Delta\phi = 0$). Secondly, the degree of observability Θ is the minimum value of \mathcal{M} . In this example, Θ is small due to $J \approx 0$ at $\Delta\phi = -9.6^\circ$ (referring back to $\hat{\phi} = 6^\circ$ and $\hat{\phi} = -3.6^\circ$) and the situation is thus poorly observable. This agrees with the qualitative discussion from Section 3.3.1.

Note that the measured quantities in J are taken as the values from the surrogate wind farm model (FLORIS) with the true ambient conditions, thus assuming a perfect model of the system. In reality, this will not hold and the work herein presents an idealized case (theoretical upper bound) of observability.

Finally, with a measure for observability defined, we can determine and analyze the observability of a particular wind farm for a certain wind direction, wind speed and turbulence intensity. The process is as follows.

1. Firstly, we evaluate the degree of observability of a single situation at a time. With a situation, we imply a wind farm subjected to a certain ambient inflow, giving us a certain set of measurements. For example, continuing the example 2-turbine wind farm of Section 3.3.1, the observability for this wind farm is investigated at a true freestream wind speed of 7.0 m s^{-1} , a freestream wind direction of 6° and a turbulence intensity of 6.5%. Referring back to Figure 3.4, our measurements would be:

$$P^{\text{measured}} = P^{\text{FLORIS}}(\phi = 6^\circ, U_\infty = 7.0 \text{ m s}^{-1}, I_\infty = 0.065) = [2.5, \quad 1.0] \cdot 10^6, \quad (3.12)$$

$$U^{\text{measured}} = U^{\text{FLORIS}}(\phi = 6^\circ, U_\infty = 7.0 \text{ m s}^{-1}, I_\infty = 0.065) = [7.0, \quad 5.5], \quad (3.13)$$

$$\phi^{\text{measured}} = [6.0, \quad 6.0] \quad (3.14)$$

The measurement vectors contain two entries, for turbine 1 and 2, respectively. In this simulation, the turbines are assumed to be aligned with the inflow wind direction; $\gamma_1 = \gamma_2 = 0$.

2. Secondly, we now assume that we do not know what the ambient conditions were that generated these measurements. This represents our estimation step. With this set of measurements, the cost function J of Equation 3.8 is calculated for a range of hypothetical (tested) ambient conditions. For this example, the estimation algorithm is limited to the estimation of U_∞ and ϕ . The (two-dimensional)

cost function is evaluated over the following ranges:

$$\Delta\phi = [-20.0 \quad -19.2 \quad -18.4 \quad \dots \quad 19.2 \quad 20.0], \text{ with } \phi = 6^\circ + \Delta\phi, \quad (3.15)$$

$$\Delta U_\infty = [-1.50 \quad -1.25 \quad -1.0 \quad \dots \quad 1.25 \quad 1.50], \text{ with } U_\infty = 7.0 \text{ m s}^{-1} + \Delta U_\infty, \quad (3.16)$$

$$\Delta I_\infty = [0.0], \text{ with } I_\infty = 0.065 + \Delta I_\infty. \quad (3.17)$$

If I_∞ is additionally to be estimated, the (three-dimensional) cost function is also evaluated over the following range for ΔI_∞ :

$$\Delta I_\infty = [-0.06 \quad -0.03 \quad 0.0 \quad 0.03 \quad 0.06], \text{ with } I_\infty = 0.065 + \Delta I_\infty. \quad (3.18)$$

Furthermore, the turbine yaw angles are fixed in the inertial frame and assumed to be known a priori in the cost function evaluations. Thus, if the cost function is evaluated for $\Delta\phi = 10^\circ$, then $\hat{\gamma}_1 = \hat{\gamma}_2 = -10^\circ$.

3. Finally, we check whether our estimation algorithm was successful. A two-dimensional (for $\Delta I_\infty = 0$) or three-dimensional (for $\Delta I_\infty = [-0.06 \quad \dots \quad 0.06]$) cost matrix is obtained following Equation 3.8, from which \mathcal{M} is calculated following Equation 3.10. The degree of observability \mathcal{O} is the minimum value of \mathcal{M} , being a positive real number.

The degree of observability \mathcal{O} can be calculated for a range of true wind directions following the process described above and displayed in a single picture. The results of such an observability analysis assuming only power measurements are available ($\lambda_P = 1$, $\lambda_U = 0$ and $\lambda_\phi = 0$) are shown in Figure 3.7 for a 6-turbine wind farm. Note that λ_U and λ_ϕ are zero to provide insight into the results. In a practical wind farm control implementation, one would opt for $\lambda_U > 0$ and $\lambda_\phi > 0$, if these measurements are available.

Each of the two radial plots shown in Figure 3.7 represents the degrees of observability for 61 different wind directions. There is one degree of observability defined for each true wind direction, plotted as a particular color across the polar axis. This thus indicates the estimability of ϕ and U_∞ for this true wind direction. For each of the 61 true wind directions, a two-dimensional cost function \mathcal{M} was calculated over the variables $\Delta\phi = [-20^\circ, -19.2^\circ, -18.4^\circ, \dots, 20.0^\circ]$ and $\Delta U_\infty = [-1.50, -1.25, -1.00, \dots, 1.50] \text{ m s}^{-1}$. Then, \mathcal{O} was taken to be the lowest value of \mathcal{M} , being the degree of observability for this true wind direction, true wind speed, true turbulence intensity, wind farm layout and with ϕ and U_∞ being the to-be-estimated parameters. We refer to this as the degree of observability for this particular ‘‘situation’’.

Figure 3.7 clearly shows that the ϕ and U_∞ can only be estimated for a narrow range of true wind directions when only power measurements are available. This makes sense, since there is only wake interaction for a small range of wind directions. Without wake interaction, one cannot distinguish, for example, between the case where all turbines operate under a yaw misalignment and a higher inflow wind speed, from the case where all turbines operate without a yaw misalignment and a lower inflow wind speed. Furthermore, an interesting difference between the observability plot for a true wind speed

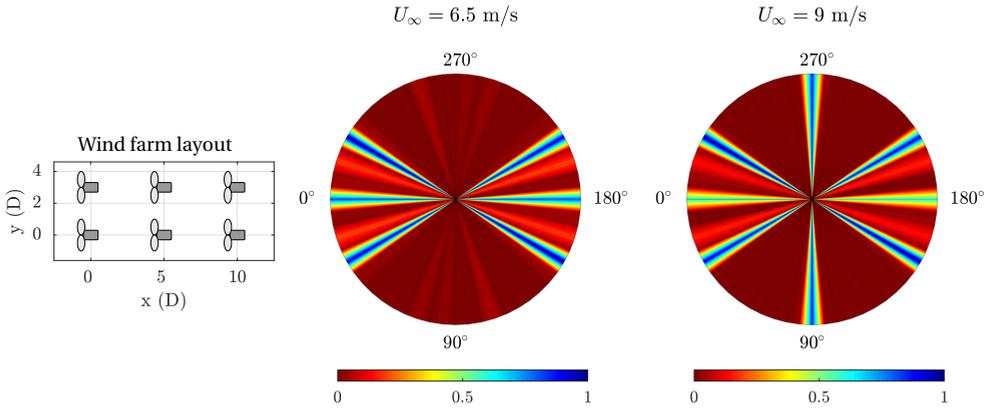


Figure 3.7: Observability plots for the 6-turbine wind farm over a range of true wind directions (from 0° to 360° , plotted in 61 discrete points along the polar axis) and two true wind speeds (6.5 m s^{-1} and 9.0 m s^{-1}). The true turbulence intensity is assumed to be known in the estimation problem, thus $J(\phi, U_\infty, I_\infty) = J(\phi, U_\infty)$ and $\Delta I_\infty = 0$, where the estimability of ϕ and U_∞ is assessed. Thus for each of the 61×2 situations (a situation is defined as a particular true wind direction and wind speed for this 6-turbine layout), the steps described earlier in this section are followed. The results are normalized to a scale of 0 to 1, with 0 being unobservable and 1 being to the best observable situation.

of 6.5 m s^{-1} and 9.0 m s^{-1} is the degree of observability at the true wind directions of 90° and 270° . This is due to the fact that the downstream turbines operate below cut-in wind speed for the 6.5 m s^{-1} case at these wind directions due to the close spacing and the wake effects. As these downstream turbines do not generate any power, their signals hold little information. For the 9.0 m s^{-1} case, all turbines operate above cut-in wind speed and thus these power signals contain more information about the flow.

The example presented in this section serves to explain the methodology and the cases become more interesting when considering more complicated farm layouts, various combinations of wind vane and wind speed measurements and the inclusion of turbulence intensity estimation. This is the focus of the next section.

3.4. A comprehensive observability analysis for 3 wind farm layouts

The observability of the ambient conditions is investigated in this section for three different wind farm layouts, namely, two symmetrical wind farms and one asymmetrical wind farm. The layouts are shown in Figure 3.8. The asymmetrical 8-turbine wind farm is an interesting configuration, as there is more unique wake interaction in this layout. This reduces the issues with symmetry previously demonstrated in Figure 3.4 compared to symmetrical wind farm layouts.

For each topology, the observability is calculated for $61 \times 4 \times 4 = 976$ situations, namely for 61 wind directions $\phi = [0^\circ, 6^\circ, 12^\circ, \dots, 354^\circ]$, 4 levels of turbulence intensity $I_\infty = [0.065, 0.095, 0.125, 0.155]$ and 4 wind speeds $U_\infty = [6.5, 9.0, 11.4, 14.5] \text{ m s}^{-1}$, of which the latter wind speed is above rated. Thus, for each of these 976 conditions, a multidimensional cost function is set-up and the most critical situation is determined following

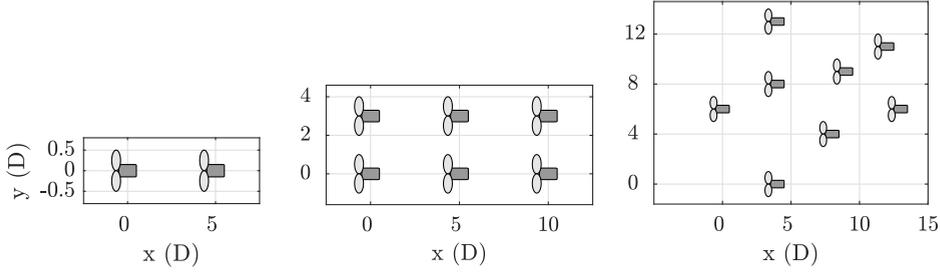


Figure 3.8: The three wind farm layouts used in a comprehensive observability analysis to demonstrate the working principles of the algorithm presented in Section 3.3. The turbines are DTU 10MW turbines (Bak et al., 2012) with a rotor diameter D of 178.3 m and a hub-height of 119 m.

Table 3.1: Relevant cost function parameters for Equation 3.9

Variable	Value
k_ϕ	$\frac{1}{40} \text{ deg}^{-2}$
k_U	$\frac{1}{3} \text{ m}^{-2} \text{ s}^2$
k_I	$\frac{1}{0.12}$
b_ϕ	4°
b_U	0.25 ms^{-1}
b_I	0.03

Equation 3.10, upon which the observability for this situation is calculated using Equation 3.9. The parameters therein are shown in Table 3.1.

This section is separated in two parts. In Section 3.4.1, the observability of the various situations is assessed under the assumption that the freestream turbulence intensity is known a priori. This simplifies the estimation problem and requires less information to be extracted from the measurements at hand. However, neglecting the estimation of I_∞ is expected to significantly worsen the accuracy of the surrogate wind farm model in a practical wind farm control algorithm. Hence, the observability with the inclusion of I_∞ is presented in Section 3.4.2.

3.4.1. Estimating ϕ and U_∞ under perfect knowledge of I_∞

First, the observability of various situations under the assumption that the turbulence intensity is known, $\hat{I}_\infty = I_\infty$, is looked into. The range over which each particular cost function is calculated is $\Delta\phi = [-20^\circ, -19.2^\circ, -18.4^\circ, \dots, 20.0^\circ]$ and $\Delta U_\infty = [-1.50, -1.25, -1.00, \dots, 1.50] \text{ m s}^{-1}$. The discretization of these parameters were tuned for convergence; such that the solutions no longer notably change at a higher precision. The range of these parameters are chosen to resemble the typical prior knowledge one has about the true ambient conditions in such an estimation problem.

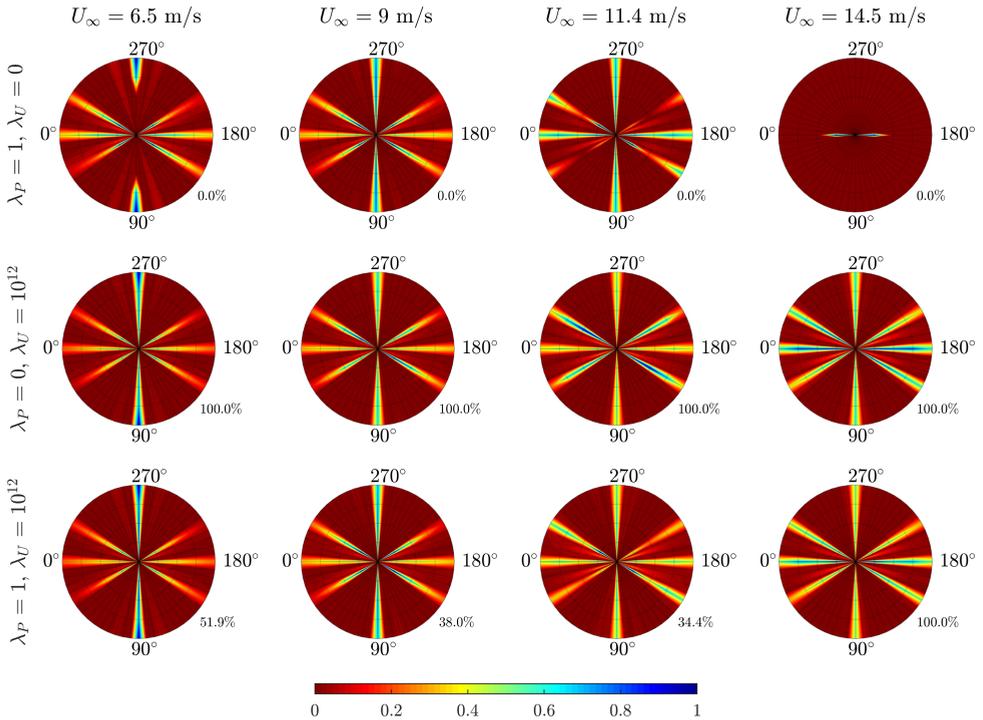


Figure 3.9: The observability for a range of wind speeds, wind directions (angular axis), turbulence intensities (radial axis) and tuning factors λ_P and λ_U under the assumption that I_∞ is known, with $\lambda_\phi = 0$. The ambient turbulence intensity varies between $I_\infty = 6.5\%$ at the center to $I_\infty = 15.5\%$ at the edges of each radial plot. The observability in each radial plot is normalized with respect to its highest value. The percentage on the bottom-right corner of each radial plot indicates to what degree the local wind speed measurements contribute to the observability. It can be seen that the power measurements provide no additional information compared to wind speed estimates and no information at all above rated wind speeds (top-right subplot).

3.4.1.1. Redundancy in the cost function: power and wind speed estimates

One important notion in the cost function shown in Equation 3.8 is that the local wind speed estimates and the turbine power signals carry duplicate information. Specifically, as the local wind speed estimators rely on the turbine power signal, the turbine power measurements theoretically add no information to the cost function that is not already included in the wind speed estimator signals. To validate this, an observability analysis is performed for the 6-turbine wind farm under $\lambda_\phi = 0$ and various values for λ_P and λ_U . The results are shown in Figure 3.9, where the dependence on turbulence intensity I_∞ is plotted along the radial axis, with $I_\infty = 6.5\%$ at the center to $I_\infty = 15.5\%$ at the edges.

From this figure, one can immediately see that situations in which all turbines are in above-rated operation are unobservable when $\lambda_U = 0$ (top-right subplot). This subplot shows some observability when the turbulence intensity is low and the wake interactions are deep, such that one or multiple downstream turbines are operating below rated

conditions. Furthermore, turbine power measurements do not add anything to the observability compared to the wind speed estimates. Note that the observability plots are not identical for below-rated conditions as power is cubically related to the wind speed, $P_i \propto U_i^3$ and thus the observability is spread slightly differently within the radial plots. Though, the trends are identical. Hence, in the remainder of this chapter, $\lambda_P = 0$.

An important remark is that a different surrogate wind farm model, e.g., one that directly correlates the upstream turbulence intensity with the upstream turbine power production, may provide a higher degree of observability from the same power measurements. Currently, such a correlation is not present in FLORIS.

3.4.1.2. Using exclusively wind speed estimator measurements ($\lambda_P = 0$, $\lambda_U = 1$, $\lambda_\phi = 0$)

Here, the situation with solely wind speed measurements available is investigated; $\lambda_P = \lambda_\phi = 0$ and $\lambda_U = 1$. This is comparable to the estimation framework applied in previous work (Doekemeijer et al., 2019b), in which wind vane measurements were not assumed to be available. This is a particularly difficult problem, as previous results from Section 3.3 suggest. In the remainder of this section, all three wind farm layouts will be addressed. The observability roses are shown in Figure 3.10.

A number of observations can be made from Figure 3.10. Firstly, for the two-turbine wind farm, it is clear that the wind direction and wind speed can only be estimated accurately for a narrow range of wind directions – specifically, in which there is sufficient wake interaction. Theoretically, the U_∞ can always be reconstructed from the wind speed estimate of the upstream turbine and the upstream turbine can be distinguished if there is wake interaction: it is the turbine with the highest power signal. The wind direction can then be estimated by looking at the quantity of wake losses at the downstream turbine. However, this may lead to situations in which two hypothesized wind directions lead to a near-identical inflow wind speed U_i , as was seen previously in Figure 3.4.

Secondly, for the six-turbine wind farm, it can be seen that this topology has more wake interaction than the two-turbine wind farm and thus has an increased observability for many situations. However, there are still situations with little to no wake interaction which are unobservable. Note that the radial plots for both the two-turbine wind farm and the six-turbine wind farm are radially symmetrical, as the topologies are also radially symmetrical.

Thirdly, for the eight-turbine wind farm, one can directly see that observability greatly increases due to many more unique wake interaction between turbines. With all topologies, generally, it is noted that a higher atmospheric turbulence leads to a lower observability. Specifically, the turbulence intensity reduces the wake interaction with downstream turbines. The results from Figure 3.10 show that ϕ and U_∞ can only be reconstructed for particular situations and thus care has to be taken in such estimation algorithms and related wind farm control algorithms. The next section shows the estimability of ϕ and U_∞ with the inclusion of wind vane measurements.

3.4.1.3. Using local wind speed and wind direction estimates ($\lambda_P = 0$, $\lambda_U = 1$, $\lambda_\phi = 10$)

By including local estimates of the wind direction, $\lambda_\phi > 0$, one can attain observability for all situations, as shown in Figure 3.11. Now, one assumes both wind speed measure-

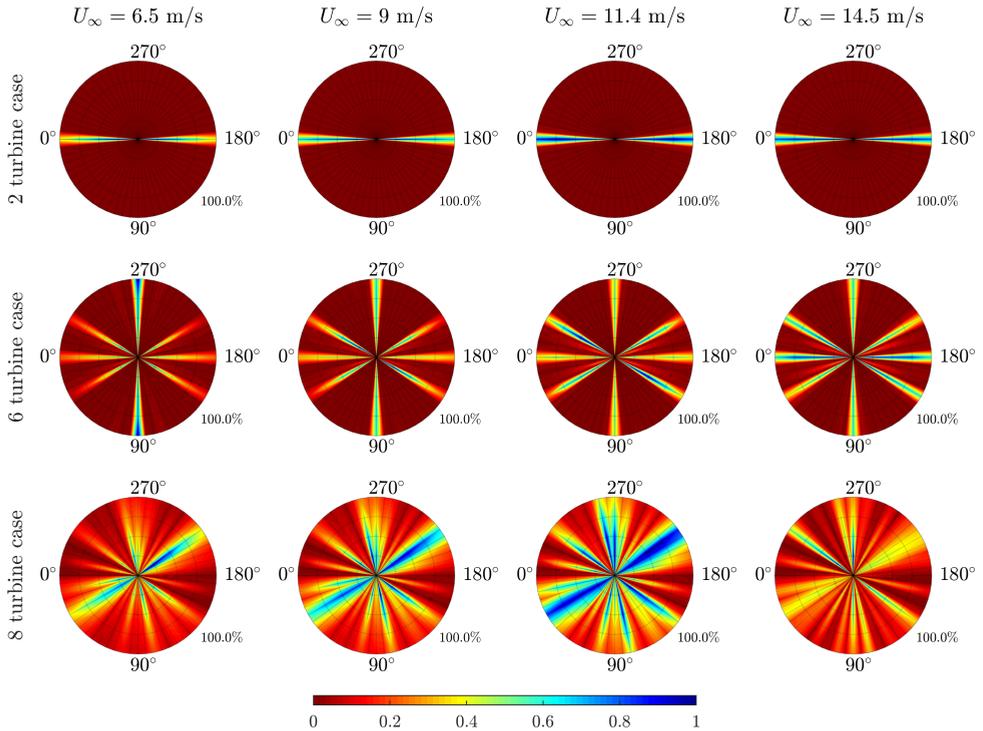


Figure 3.10: The observability for a range of wind speeds, wind directions (angular axis), turbulence intensities (radial axis) and wind farm layouts under the assumption that I_∞ is known, with $\lambda_P = \lambda_\phi = 0$. The ambient turbulence intensity varies between $I_\infty = 6.5\%$ at the center to $I_\infty = 15.5\%$ at the edges of each radial plot. The observability in each radial plot is normalized with respect to its highest value. The percentage on the bottom-right corner of each radial plot indicates to what degree the local wind speed measurements contribute to the observability, which in this situation is 100%.

ments and wind vane measurements to be available.

It is clear to see that all the necessary information is contained in the measurements available for the estimation of U_∞ and ϕ : all situations appear observable. Observability is guaranteed due to the availability of local wind speed and wind direction measurements, which are quantities directly derived from the ambient wind speed, ambient wind direction and the wake interactions. Note that there are some variations within the radial circle, which are both due to physical effects such as more or less wake interaction and also due to fact that the search space of the cost function ($\Delta\phi$, ΔU_∞ , ΔI_∞) is discretized at a finite resolution.

The tools presented in this chapter may prove useful to find a balanced trade-off in the cost function between the contributions from various measurement sources. However, even with an accurate estimation of ϕ and U_∞ , significant model discrepancies may remain. The freestream turbulence intensity I_∞ has a relatively large impact on the optimal turbine setpoints for wake steering, as it has a direct relationship to the degree of

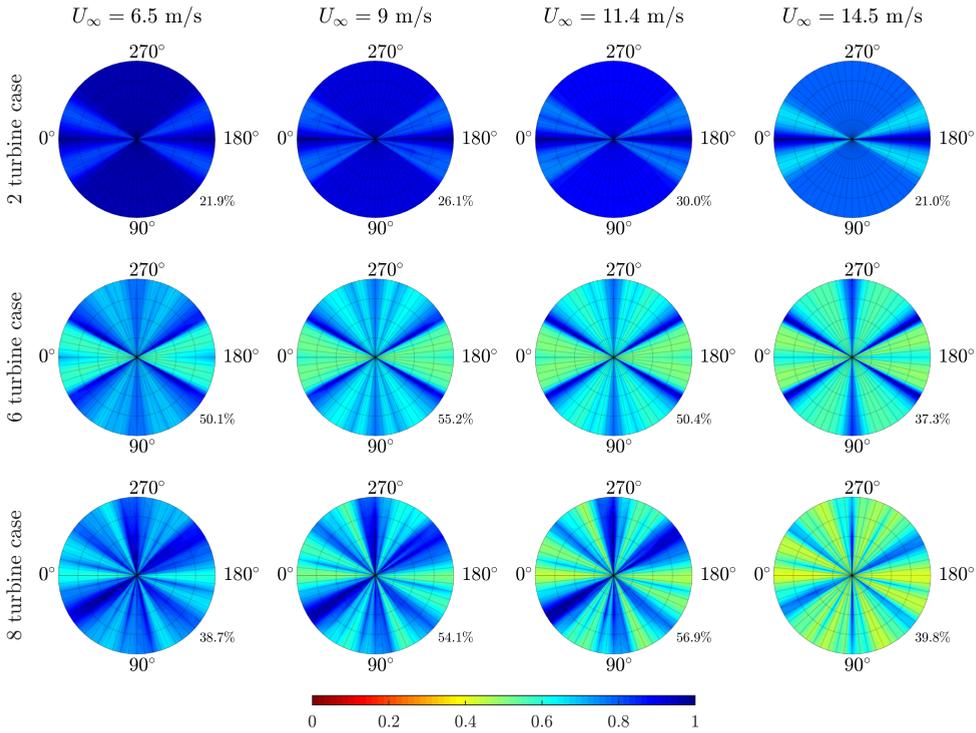


Figure 3.11: The observability for a range of wind speeds, wind directions (angular axis), turbulence intensities (radial axis) and wind farm layouts under the assumption that I_∞ is known, with $\lambda_U = 1$ and $\lambda_\phi = 10$. The ambient turbulence intensity varies between $I_\infty = 6.5\%$ at the center to $I_\infty = 15.5\%$ at the edges of each radial plot. The observability in each radial plot is normalized with respect to its highest value. The percentage on the bottom-right corner of each radial plot indicates to what degree the local wind speed measurements contribute to the observability, which provides an idea to the robustness of the solution.

wake recovery. Hence, the estimation of I_∞ is a necessity in reliable wind farm control algorithms. In the next section, the estimation of I_∞ is incorporated into the observability analysis.

3.4.2. The full estimation problem: estimating ϕ , U_∞ and I_∞

While observability for all situations was shown in Section 3.4.1.3, a compromising assumption was made that the freestream turbulence intensity I_∞ was known. In reality, this is not a realistic assumption and I_∞ must be estimated together with U_∞ and ϕ . The observability when estimating ϕ , U_∞ and I_∞ is shown in Figure 3.12, where $\Delta I_\infty = [-0.06, -0.03, 0.0, 0.03, 0.06]$.

Several observations can be made. Firstly, one can directly see that the observability significantly reduces for a large range of conditions compared to only the estimation of ϕ and U_∞ . For the two-turbine case, observability only remains for the narrow window of wind directions in which there is wake interaction. This can be explained by the fact

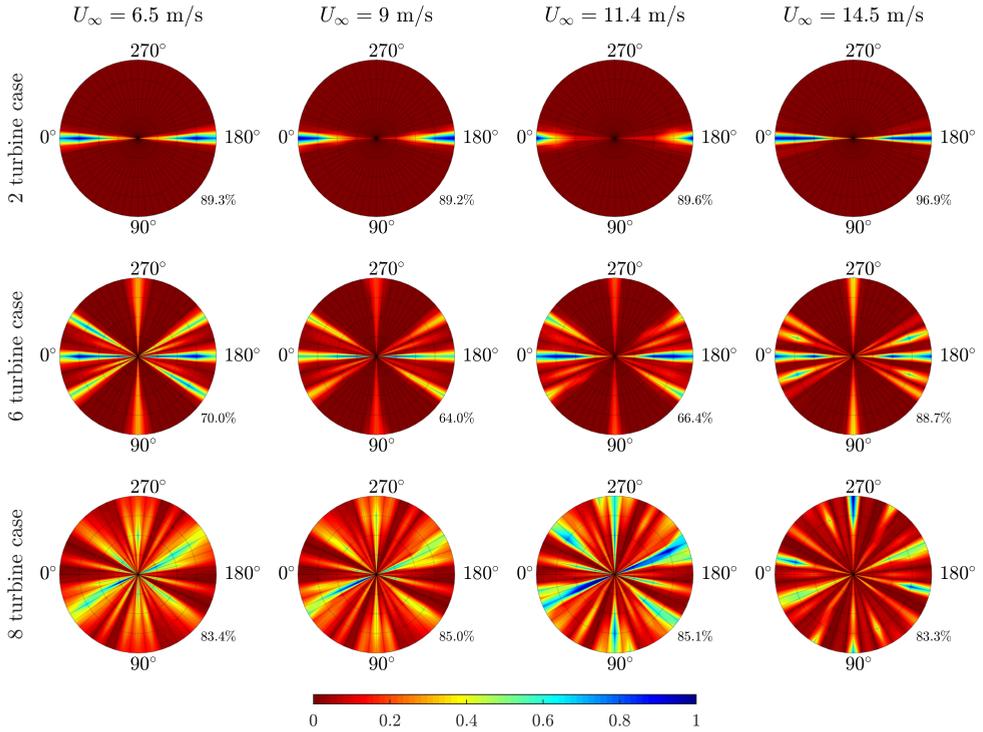


Figure 3.12: The observability for a range of wind speeds, wind directions (angular axis), turbulence intensities (radial axis) and wind farm layouts with $\lambda_T = 1$ and $\lambda_\phi = 10$. The ambient turbulence intensity varies between $I_\infty = 6.5\%$ at the center to $I_\infty = 15.5\%$ at the edges of each radial plot. The observability in each radial plot is normalized with respect to its highest value. The percentage on the bottom-right corner of each radial plot indicates to what degree the local wind speed measurements contribute to the observability, which provides an idea to the robustness of the solution.

that the measurements provide direct information on ϕ and U_∞ , while the estimation of I_∞ is enabled through inversion of the surrogate wind farm model and the usage of the local wind speed measurement at the downstream turbine. This only applies when wake interaction is present.

Secondly, observability is reduced in the 6-turbine case compared to Figure 3.11, yet observability remains more widespread than the two-turbine case. More wake interaction and multiple-wake interaction leads to the fact that the turbine power signals are more sensitive to the freestream turbulence and thus yield a higher observability than the two-turbine case. Additionally, while a higher turbulence intensity leads to additional wake recovery, it also leads to wider wakes which can impact a downstream turbine where it would not for lower turbulence intensities. These two effects have an opposite effect on the observability and hence observability does not uniformly decrease with an increase in the freestream turbulence intensity.

Thirdly, the 8-turbine wind farm has the most observable situations from the three

topologies. Due to the many unique wake interactions, the solutions become relatively sensitive to the freestream turbulence intensity and the ambient conditions can be estimated for most conditions. Though, also in this wind farm one can find several situations in which the freestream conditions cannot uniquely be reconstructed from the measurements available.

An important remark to make is that all results presented in this section ignore the possibility of other measurement sources. While this framework allows the inclusion of turbulence intensity measurements, this was not pursued here. Additionally, one may argue that temporal correlation of measurements would allow for additional information on the ambient conditions. This would require a dynamic mathematical model that correlates the ambient conditions and the turbine measurements and a state estimation algorithm such as a Kalman filter. This is out of the scope of this thesis.

Finally, recall that these results present an idealized case, in which there is no measurement noise and the surrogate wind farm model is used to generate the measurements, implying that the surrogate model perfectly represents reality. None of these assumptions are valid in practice and thus the observability roses presented in this section will further diminish. Though, the results presented in this section are an useful step towards the synthesis of an algorithm that estimates the ambient conditions in a robust manner. The observability roses from Figure 3.12 provide a theoretical upper limit on the relative estimatibility of the ambient conditions ϕ , U_∞ and I_∞ from the measurements available. This can provide guidance in wind farm control algorithms on when to estimate certain parameters. Since ϕ and U_∞ are always estimable according to Figure 3.11, the observability analysis presented in this section can be used to determine whether to estimate I_∞ in addition to U_∞ and ϕ . If the situation is “observable enough” (which is to be selected experimentally), the measurements should contain sufficient information to reliably estimate I_∞ . If not, one can assume I_∞ to be equal to its past value (since the turbulence intensity also does not change very rapidly in the field) and exclusively estimate ϕ and U_∞ . This approach is currently being explored and will be published in future work.

3.5. Conclusions

Over the last years, the scientific community surrounding wind farm control has shown an increasing amount of interest towards the real-time estimation of the ambient conditions inside a wind farm. This ambient flow information is essential to the optimization of the turbine yaw angles for wake steering, which is currently the most popular methodology of wind farm control for power maximization. The degree of reconstructability of the ambient conditions highly depends on the measurements available and the wind farm layout. For many situations, it is clear to see that the ambient conditions cannot be estimated. However, no quantitative measure exists to represent the degree of estimability of the ambient conditions. This chapter addressed this scientific gap.

The main contribution of this chapter was the introduction of a novel, mathematical definition for the observability of the ambient conditions. This measure describes how well the true ambient conditions can be distinguished from hypothesized ambient con-

ditions through a model inversion approach for a particular set of measurements. This measure of “observability” is modular and can easily be extended with other measurement sources or other surrogate models. While a number of outcomes of this chapter may seem apparent, this theoretical framework provides the tools for extended analysis and a quantitative measure for the estimability of the inflow properties leveraging different measurement sources and surrogate wind farm models.

In several case studies, we showed the usefulness of the proposed measure. Moreover, while information concerning the wind direction can be derived by looking at the turbine power signals and the inter-turbine wake interactions, the presented case studies showed that there is a strong need for wind direction measurements for WFC to sufficiently cover observability for any topology and any ambient condition. Generally, situations in which there is sufficient wake interaction are observable, while situations with little to no wake interactions are unobservable.³ Also, more complicated, unstructured wind farm layouts generally result in a higher observability as there are more unique wake interactions between turbines. Furthermore, the mathematical framework supports the straight-forward notion that local turbine power measurements provide no additional information compared to local wind speed estimates, implying that power measurements can be omitted from the cost function.⁴

In general, even with local wind speed and wind direction information, one still cannot reconstruct the full set of ambient conditions (wind speed, wind direction and turbulence intensity) for all conditions that a particular wind farm may encounter. Thus, before one may attempt to estimate the ambient conditions, one should consider whether the situation is observable in the first place. Using this information, one may condition their wind farm control algorithm to situations that are sufficiently observable. This aids in improving the reliability of wind farm control algorithms and thereby hopefully the willingness to adopt such algorithms by the industry.

³Here, the availability of other measurements than the local wind direction, wind speed and turbine power capture were neglected. Additionally, the temporal evolution of measurement signals may provide insight into the current ambient conditions. This was outside of the scope of this work.

⁴A promising research direction is described by Schreiber et al. (2020), in which model errors are corrected for by superimposing a correctional model based on operational data. This concept is not further explored here.

4

STEADY-STATE-MODEL-BASED WIND FARM CONTROL UNDER TIME-VARYING INFLOW

ABSTRACT

Wind farm (WF) controllers adjust the control settings of individual turbines to enhance the total performance of a wind farm. Most WF controllers proposed in the literature assume a time-invariant inflow, whereas important quantities such as the wind direction and speed continuously change over time in reality. Furthermore, properties of the inflow are often assumed known, which is a fundamentally compromising assumption to make. This chapter presents a novel, closed-loop WF controller that continuously estimates the inflow and maximizes the energy yield of the farm through yaw-based wake steering. The controller is tested in a high-fidelity simulation of a 6-turbine wind farm. The WF controller is stress-tested by subjecting it to strongly time-varying inflow conditions over 5000 s of simulation. A time-averaged improvement in energy yield of 1.4% is achieved compared to a baseline, locally greedy wind farm controller. Moreover, the instantaneous energy gain is up to 11% for wake-loss-heavy situations. Note that this is the first closed-loop and model-based WF controller tested for time-varying inflow conditions (i.e., where the mean wind direction and wind speed change over time) at such fidelity. This solidifies the WF controller as the first realistic closed-loop control solution for yaw-based wake steering.

The contents of this chapter have been published as a separate research article in the peer-reviewed Renewable Energy journal, as:

B M Doekemeijer, D C van der Hoek, and J W van Wingerden. Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions. *Renewable Energy*, 156:719–730, 2020b. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2020.04.007>.

4.1. Introduction

In the “Klimaatakkoord” (Rijksoverheid, 2019), the Dutch government pledges to significantly reduce carbon-dioxide emissions over the next decade, lowering emissions to 49% of the levels measured in 1990. In pursuit of this goal, the Dutch government has been installing an increasing number of wind farms in the North Sea. The intention is to have 11 GW of wind energy installed off the coast of The Netherlands by 2030, thereby accounting for 40% of the Dutch national electricity demand.

As the globally installed capacity of wind energy continues to grow, so does the interest towards further improving the efficiency of wind turbines and wind farms. The research field of control engineering plays a significant role in this process. In the past, the focus of control engineering has been on individual wind turbine control. More recently, the focus has shifted from wind turbine towards wind farm control, in which turbines are coordinated with one another to achieve a collective objective (Boersma et al., 2017).

A distinction can be made between two strategies of wind farm control, namely *induction control* and *wake steering*. The first strategy aims at derating upstream turbines, which purposely lowers their energy yield through pitching the turbine blades and adjusting the generator torque. Induction control is typically used for active power control (e.g., Boersma et al., 2019a; Shapiro et al., 2018), load mitigation strategies (e.g., Kanev et al., 2018) and more recently power maximization (e.g., Munters and Meyers, 2017). The second strategy, wake steering, displaces the wake downstream by purposely misaligning the rotor plane with the incoming air stream. Wake steering is typically done using yaw control. The most common objective of wake steering is power maximization (e.g., Fleming et al., 2019; Gebraad et al., 2016; Howland et al., 2019).

Yaw-based wake steering has shown significant potential in high-fidelity simulations and real-world experiments. For example, Gebraad et al. (2016) show an increase in energy extraction of up to 13% for a 6-turbine wind farm in a large-eddy simulation. Furthermore, Campagnolo et al. (2016a) show an increase in energy extraction of up to 21% for an array of 3 turbines in a wind tunnel. Additionally, Fleming et al. (2019) show an increase in energy extraction of up to 4% for an array of 3 turbines in a field experiment. Moreover, Howland et al. (2019) demonstrate wake steering through field experiments on an array of 6 wind turbines, showing an increase of up to 47% at low wind speeds (due to cut-in behaviour of downstream turbines) and an increase of up to 13% for higher wind speeds. Due to its promising potential, the chapter at hand focuses on yaw-based wake steering for power maximization.

However, a crucial limitation of most wind farm controllers proposed in the literature is that the incoming wind field is assumed to be time invariant. In reality, the wind field entering a wind farm changes continuously due to fluctuations in the heating of the Earth's surface, among others. Moreover, wake steering is highly sensitive to the ambient conditions (Boersma et al., 2017). In consequence, it is crucial to test farm control solutions for realistic, time-varying inflow conditions. Additionally, properties of the inflow such as the wind direction and wind speed are typically assumed to be prior knowledge to the wind farm controller, which is a secondary unrealistic and compromising assumption. Wind farm controllers that do not rely on live measurements but rather on

prior knowledge of the inflow are denoted as *open-loop*. The results from Gebraad et al. (2016) and Campagnolo et al. (2016a) are with open-loop wind farm controllers under time-invariant inflows, in which the inflow is carefully selected to maximize the potential benefits of wake steering.

Furthermore, the energy yield gain of up to 4% shown by Fleming et al. (2019) is extracted from months of field data, for narrow time windows and inflow conditions. This provides little insight into the benefits of wake steering over the annual operation cycle of a wind farm. Furthermore, despite the notable improvements shown in Howland et al. (2019), the gain in annual energy production for the field experiments was found to be insignificant.

The wind farm controllers deployed in both field experiments (Fleming et al., 2019; Howland et al., 2019) are open-loop, deriving the inflow properties from the most upstream turbine or an external measurement system. However, the uncertainties concerning inflow estimation and the high complexity in modeling the relevant wind farm dynamics require a *closed-loop* wind farm control solution. In closed-loop control, measurements of the controlled system are fed back to the controller to allow adaptation to a changing environment and model uncertainty. The closed-loop model-based wind farm control framework presented in this chapter consists of two components, being *model adaptation* and *setpoint optimization*, as depicted in Figure 4.1. Model adaptation consists of estimating the input parameters of a surrogate model that are currently relevant for the wind farm. This surrogate model is a simplified mathematical model of the wind farm dynamics with a low computational cost. In practice and also in this chapter, model adaptation often implies the estimation of the freestream wind speed, wind direction and the amount of wake recovery (commonly defined by the turbulence intensity). Secondly, the setpoint optimization leverages the adapted surrogate model to find the turbine control setpoints that maximize a certain objective. In this chapter, the objective is power maximization and the control variables are the turbine yaw angles.

The main focus in the wind farm control literature has been on surrogate model development (e.g., Bastankhah and Porté-Agel, 2016; Gebraad et al., 2016) and control setpoint optimization (e.g., Annoni et al., 2018a; Bay et al., 2018; Marden et al., 2013). More recently, there has been an increasing amount of interest towards the estimation of the ambient conditions and dealing with the time-varying nature of wind (e.g., Chapter 3 of this dissertation and Annoni et al., 2019; Rott et al., 2018; Simley et al., 2020). However, many of the wind farm control algorithms proposed in the literature are merely tested on simplified simulation models, from which no real conclusions can be drawn apart from a proof of concept. Furthermore, the controllers that are tested in high-fidelity (large-eddy) simulations and real-world experiments typically assume a constant mean inflow wind direction, wind speed and turbulence intensity (e.g., Campagnolo et al., 2016a; Gebraad et al., 2016). However, experiments under such steady inflow conditions insufficiently represent real-world scenarios and thus much uncertainty remains concerning the true potential of these wind farm controllers in actual farms.

A handful of articles exist that consider time-varying inflow conditions for wind farm controller validation. Bossanyi (2019) demonstrates a wind farm control algorithm in

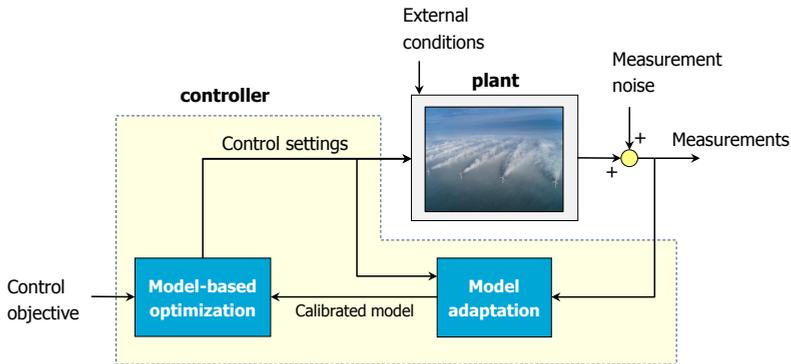


Figure 4.1: The architecture for closed-loop model-based wind farm control. This architecture combines the benefits of closed-loop model-free with open-loop model-based control, being quick convergence rates and resilience to model uncertainties, respectively. Here, turbines are controlled according to a simplified mathematical model of the wind farm. Additionally, the effects of the applied control policy is measured and taken into account in deciding the next control policy, either directly or through recalibration of the surrogate wind farm model (exemplified here).

low-fidelity simulation subjected to time-varying inflow conditions. Also, Vollmer (2018) demonstrates open-loop wake steering on a two-turbine array in a large-eddy simulation subjected to time-varying inflow conditions. Furthermore, Ciri et al. (2017b) presents a closed-loop and model-free wind farm control algorithm that improves the performance of turbines inside a wind farm, demonstrated in high-fidelity simulations under a time-varying inflow. However, model-free algorithms for wake steering are fundamentally limited due to slow convergence rates combined with the inherent variability of the inflow conditions. Their practicability in real wind farms therefore remains uncertain (Boersma et al., 2017).

To the best of the authors' knowledge, there is no literature on the assessment of closed-loop model-based wind farm control solutions in a high-fidelity environment (i.e., field experiment, wind tunnel experiment, large-eddy simulation) with time-varying inflow conditions. Addressing this scientific gap is invaluable for the practical validation and implementation of wind farm control solutions, as time-varying inflow conditions are ubiquitous in real-world wind farms. This chapter contains three novel contributions:

1. a detailed fit of the surrogate wind farm model FLORIS to large-eddy simulation data.
2. a model-based estimation algorithm that predicts the freestream wind direction, wind speed and a wake recovery factor using measurements that are readily available in commercial wind farms. This algorithm is assisted by a theoretical measure of observability presented in Chapter 3 of this dissertation.
3. validation of the closed-loop wind farm controller in a large-eddy simulation sub-

jected to time-varying inflow conditions.

The structure of this chapter is as follows. In Section 4.2, the high-fidelity simulation environment is described. In Section 4.3, the surrogate wind farm model is outlined. In Section 4.4, the closed-loop wind farm controller is synthesized. This controller is tested in Section 4.5 in a high-fidelity simulation under time-varying inflow conditions, upon which the turbine energy yield and the turbine loads are investigated. The chapter is concluded in Section 4.6.

4.2. The Simulator for Wind Farm Applications

For surrogate model tuning and controller validation, the high-fidelity Simulator for Wind Farm Applications (SOWFA) model developed by the National Renewable Energy Laboratory (NREL) is used in this chapter. SOWFA is a large-eddy wind farm simulation model that leverages the actuator line model to determine the forces applied by each turbine on the flow. SOWFA has been validated to SCADA data in Churchfield et al. (2012b), among others. In recent work, this high-fidelity simulator was coupled with MATLAB using a network-based communication interface (Doekemeijer et al., 2019b) for straight-forward control algorithm testing. In this chapter, the wind turbine of interest is the DTU 10MW reference wind turbine (Bak et al., 2012). An overview of important parameters for all simulations presented in this chapter is given in Table 4.1. The reader is referred to Churchfield et al. (2012a) for a more detailed description of SOWFA.

4.3. Surrogate model

The closed-loop control architecture outlined in Figure 4.1 requires a surrogate model, serving for both the model adaptation and the control setpoint optimization. Therefore, this model should predict the flow and turbine behavior in the farm accurately while being computationally tractable for real-time application. For this purpose, a popular, steady-state, control-oriented surrogate wind farm model is opted for. To increase accuracy, several model parameters are calibrated in accordance to high-fidelity simulation data.

4.3.1. Model definition

The surrogate model employed in this chapter is the FLOW Redirection and Induction in Steady-state (FLORIS) model (Doekemeijer and Storm, 2019). FLORIS predicts the time-averaged three-dimensional flow field and turbine power capture of a wind farm as a function of the turbine control settings and the incoming wind field. Since the surrogate model is static, the computational cost for a single model evaluation is between 10 milliseconds and 1 second. This low computational cost makes FLORIS feasible for real-time control applications. The general inputs and outputs of the FLORIS model are outlined in Figure 4.2, categorized into four input classes (*ambient conditions*, *control settings*, *wind farm properties* and *model definition*) and two output classes (*turbine outputs* and *flow field*). For a more detailed description of FLORIS, the reader is referred to the literature (Bastankhah and Porté-Agel, 2016; Crespo and Hernández, 1996; Katic et al., 1987).

Table 4.1: Important variables for the large-eddy simulations in this chapter

Variable	Value
Turbine type	DTU 10MW (Bak et al., 2012)
Hub height	119.0 m
Rotor diameter (D)	178.3 m
Turb. spacing stream-wise (if appl.)	5.0 D
Turb. spacing cross-stream (if appl.)	3.0 D
Rotor approximation	Actuator line (<i>ALMAdvanced</i>)
Domain size	16.8 $D \times 16.8 D \times 5.6 D$
Cell size (base mesh)	10 m \times 10 m \times 10 m
Cell size (refined, near rotor)	2.5 m \times 2.5 m \times 2.5 m
Blade epsilon	5.0 m (Troldborg, 2008)
ABL stability	Neutral
Inflow wind speed (U_∞): case A	7.0 m/s
Inflow wind speed (U_∞): case B	8.2 m/s
Inflow wind speed (U_∞): case C	7.3 m/s
Inflow wind speed (U_∞): case D	Time-varying, 6.0 – 10.0 m/s
Inflow wind direction (ϕ): case A	0° (west)
Inflow wind direction (ϕ): case B	0° (west)
Inflow wind direction (ϕ): case C	0° (west)
Inflow wind direction (ϕ): case D	Time-varying, 0° – 90° (south-west)
Surface roughness (z_0): case A	0.0 m ($I_\infty = 0\%$)
Surface roughness (z_0): case B	$2.0 \cdot 10^{-4}$ m ($I_\infty = 5\%$)
Surface roughness (z_0): case C	2.0 m ($I_\infty = 12\%$)
Surface roughness (z_0): case D	$2.0 \cdot 10^{-4}$ m ($I_\infty = 7 - 13\%$)
Time step	0.20 s

Model discrepancies between SOWFA and FLORIS are inevitable. In FLORIS, several assumptions are made on the single wake profile and the interaction between multiple wakes. In addition to the absence of secondary steering effects (Martínez-Tossas et al., 2019), the lack of time-dependent wake propagation in FLORIS is a significant source of model discrepancies.

4.3.2. Model tuning prior to controller synthesis

Surrogate wind farm models typically include a myriad of model parameters (e.g., Bastankhah and Porté-Agel, 2016; Gebraad et al., 2016). FLORIS has 10 free parameters that must be defined prior to controller synthesis. Typically, the values of these parameters are based on idealized theory and wind tunnel experiments (Bastankhah and Porté-Agel, 2016). Moreover, these parameters are known to vary with wind turbine type and various wind farm properties (Doekemeijer et al., 2019b). The success of the controller largely relies on the accuracy of the surrogate model. Hence, in this chapter, the model parameters are tuned prior to controller synthesis in accordance to high-fidelity simulation data

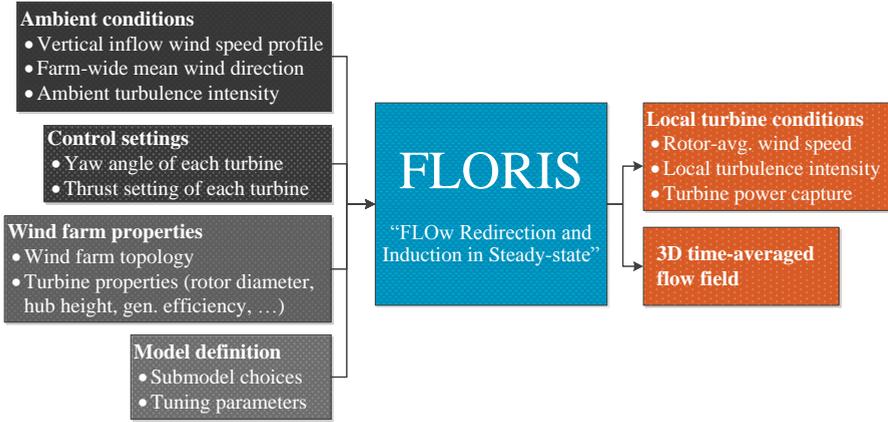


Figure 4.2: Flowchart of the FLORIS model. This model has four classes of inputs: the ambient conditions, the turbine control settings, the wind farm properties (e.g., layout) and a set of model parameters. FLORIS maps these inputs in a static fashion to a set of turbine outputs being the power capture and the three-dimensional flow field.

of the wind turbine and wind farm of interest.

Firstly, the power curve of a single turbine as a function of the yaw misalignment is tuned using cases A and B of Table 4.1. Currently, FLORIS relies on a database of power and thrust coefficients, C_P and C_T , for the DTU 10MW turbine. This database includes the effect of yaw on the power production and was generated using blade element momentum (BEM) theory. Due to a difference in BEM theory and large-eddy simulation, the power coefficient database is scaled by an empirically found multiplication factor of $\eta(\gamma) = \frac{1.08}{\cos \gamma}$. The result is shown in Figure 4.3.

Secondly, the wind profile behind a single turbine is tuned. This is done by minimizing the error in the predicted flow fields between FLORIS and SOWFA for cases A, B and C of Table 4.1 using a genetic algorithm. For each case, 7 yaw setpoints are assessed: from $\gamma = -30^\circ$ to $\gamma = 30^\circ$ in steps of 10° . The optimal model parameters Ω^* are found by minimizing the root-mean-square error (RMSE) of the time-averaged flow field from SOWFA, $U^{\text{SOWFA}} \in \mathbb{R}^{N_u}$ and the flow field predicted by FLORIS, $U^{\text{FLORIS}} \in \mathbb{R}^{N_u}$, as

$$J_{\text{fit}}(\Omega) = \frac{1}{N_u} \sum_{i=1}^{N_u} w_i (U_i^{\text{FLORIS}}(\Omega) - U_i^{\text{SOWFA}})^2. \quad (4.1)$$

Here, the U -vectors are populated by taking N_u samples from the vertical cross-stream slices (wake profiles) at $x = 3D$, $x = 5D$, $x = 7D$ and $x = 10D$ downstream, with D the rotor diameter. Furthermore, the weighing factors are chosen as $w_i^{3D} = 1$, $w_i^{5D} = 2$, $w_i^{7D} = 2$ and $w_i^{10D} = 1$, respectively, to emphasize a good calibration in the flow field at $5D$ and $7D$ downstream, as turbines are often sited at this distance. Table 4.2 shows the optimized model parameters Ω^* and the lower and upper optimization constraints.

Interesting to note is that, since FLORIS is now tuned to time-averaged SOWFA data,

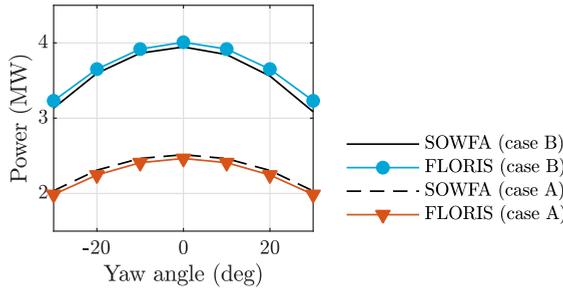


Figure 4.3: The power curve according to FLORIS and according to SOWFA as a function of the yaw angle. FLORIS is tuned to the data of two SOWFA simulations (cases A and B of Table 4.1) to guarantee a better match in power capture due to a yaw misalignment.

Table 4.2: Optimal model parameters Ω^* after model fitting, including optimization bounds (*min* and *max*).

Symbol	Relates to the...	Min	Max	Value
k_a	wake expansion (Bastankhah and Porté-Agel, 2016)	0.05	1.5	0.537
k_b	wake expansion (Bastankhah and Porté-Agel, 2016)	-0.01	0.02	-0.00085
a_d	wake deflection (Bastankhah and Porté-Agel, 2016)	-1.0	1.0	0.0011
b_d	wake deflection (Bastankhah and Porté-Agel, 2016)	-0.1	0.1	-0.0077
α	near-wake region (Bastankhah and Porté-Agel, 2016)	0.5	10.0	1.088
β	near-wake region (Bastankhah and Porté-Agel, 2016)	0.03	0.60	0.222
τ_a	induced turbulence (Crespo and Hernández, 1996)	0.07	10.0	7.84
τ_b	induced turbulence (Crespo and Hernández, 1996)	0.08	10.0	4.57
τ_c	induced turbulence (Crespo and Hernández, 1996)	0.001	0.50	0.43
τ_d	induced turbulence (Crespo and Hernández, 1996)	-5.0	-0.01	-0.246

FLORIS implicitly includes the time-averaged impact of wake meandering on the flow and on the power production of turbines.

4.3.3. Model validation

The calibrated surrogate wind farm model is validated using a different, unused dataset. The optimal parameter set Ω^* is compared to three simulations of case B (Table 4.1) and a three-turbine wind farm spaced $5D$ apart. In these simulations, the first two turbines are misaligned with the inflow at $\gamma_{1,2} = -20^\circ$, $\gamma_{1,2} = 0^\circ$ and $\gamma_{1,2} = 20^\circ$, respectively, with $\gamma_3 = 0^\circ$ for all three cases. This layout, inflow and operating conditions are chosen for a number of reasons. Firstly, the second turbine is set up to experience a slower, more turbulent inflow than the upstream turbine, effectively testing the wake model for inflow conditions it was not tuned for. Secondly, the third turbine operates in partially waked inflow, which is a common condition often causing significant model discrepancies in surrogate models (Martínez-Tossas et al., 2019). Thirdly, the upstream two turbines are purposely yawed in either direction to assess the model's validity under realistic wake steering.

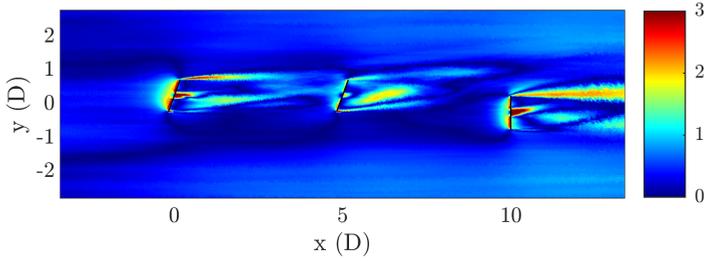


Figure 4.4: The absolute error between the hub-height flow field from SOWFA and as predicted by FLORIS in m/s. A good match is seen in the far-wake regions and in front of the downstream turbines. While the near-wake region is difficult to predict, this region is less important for wind farm control.

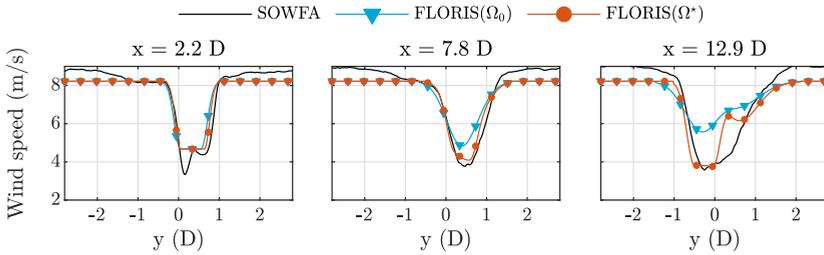


Figure 4.5: The wake profile at hub height for different locations downstream. An improvement is seen for the optimized set of parameters, Ω^* , compared with the default parameters, Ω_0 , from the literature (Bastankhah and Porté-Agel, 2016; Crespo and Hernández, 1996; Katic et al., 1987). Note that the wind speed outside of the wake appears higher in SOWFA. This is due to non-homogeneous effects in the turbulent inflow modeled in SOWFA.

Figure 4.4 shows the absolute error between the hub-height flow field from SOWFA and as predicted by FLORIS for one validation case. This figure clearly shows that most errors are in the near-wake region, which are not of interest for wind farm control. Generally, the far-wake regions and the flow in front of downstream turbines are well predicted, which should in turn lead to accurate predictions of the energy yield.

Furthermore, Figure 4.5 shows the cross-sectional wake profile at the turbine hub height at several positions downstream, $x = 2.2 D$, $x = 7.8 D$ and $x = 12.9 D$. Generally, the conclusions drawn from Figure 4.4 are confirmed. Additionally, this figure clearly shows an improvement of the parameter set Ω^* over the default parameter choice Ω^0 .

In conclusion, a good match is found between FLORIS with Ω^* and the time-averaged results from SOWFA for unseen data with multiple turbines and more complicated wake interaction. This analysis brings sufficient confidence in the surrogate model for it to be used in controller synthesis.

4.3.4. Introducing a wake recovery factor

While FLORIS has often shown a good match with high-fidelity data (e.g., Doekemeijer et al., 2019b; Gebraad et al., 2016) and experimental field data (Annoni et al., 2018b), the

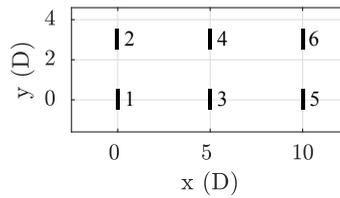


Figure 4.6: The 6-turbine wind farm considered in this chapter. The wind farm has a spacing of $5 D$ longitudinally and $3 D$ laterally.

amount of wake recovery in FLORIS is described by a single input variable, being the freestream turbulence intensity I_∞ . However, simulations show that the model accuracy of FLORIS can improve by assigning a turbulence intensity different from the true physical value – leading to a better prediction of the flow field and the turbine power signals. Therefore, the turbulence intensity I_∞ is hereafter redefined as the *wake recovery factor* ζ . This parameter can be interpreted as a proxy for the amount of wake recovery, with $\zeta = 0.0$ corresponding to a situation without wake recovery and $\zeta = 1.0$ corresponding to a situation with maximum wake recovery. Section 4.4.1 will present how the wind farm controller calculates ζ using measurements in the wind farm.

4.4. Controller synthesis

The surrogate wind farm model of Section 4.3 is used to design a closed-loop wind farm controller. The wind farm studied in this chapter is a virtual offshore wind farm with six DTU 10MW turbines (Bak et al., 2012) spaced at $5 D \times 3 D$ as shown in Figure 4.6. The model adaptation algorithm is described in Section 4.4.1. The control setpoint optimization algorithm is described in Section 4.4.2. An overview of the controller is given in Section 4.4.3.

4.4.1. Real-time model adaptation

Performance of the controller is highly dependent on the assumed wind direction, wind speed and the amount of wake recovery inside the surrogate model. As not all of these variables are measured accurately in the farm, a wind-farm-wide estimation must be made before the control setpoints are optimized.

In previous work (Doekemeijer et al., 2019b), the wind direction was assumed to be estimated using the approach of Bertelè et al. (2017), through the usage of blade load measurements and BEM theory to derive the turbine inflow conditions. However, blade load sensors are typically not available in commercial wind turbines. In this chapter, rather, a temporally and spatially averaged freestream wind speed, wind direction and wake recovery factor are estimated using the readily available generator power and wind direction measurements of each turbine. Thus, this control solution does not require additional sensors to be installed in the wind farm. Moreover, previous work did not consider the (lack of) observability of the ambient conditions for particular situations. The observability measure presented in Chapter 3 of this dissertation is now included

in the algorithm to decide which parameters can be estimated from the measurements available. The complete estimation algorithm follows a sequential approach:

1. The freestream wind direction ϕ is estimated by turbine- and time-averaging the local turbines' estimates of the wind direction. In commercial turbines, such estimators readily include low-pass filtering, bias and drift correction. For simplicity, the wind direction measurements are idealized and taken as the wind direction setpoints assigned to SOWFA in this chapter. To increase realism, these measurements are disturbed by artificial Gaussian noise representing measurement noise. Further, the turbine yaw angles γ in FLORIS are based on the nacelle orientation measurement averaged over a set time horizon.
2. The freestream wind speed is estimated from the upstream turbines. The set containing the indices of upstream turbines is denoted by \mathcal{U} , with $N_{\mathcal{U}}$ the number of upstream turbines. Mathematically, we solve

$$U_{\infty} = \operatorname{argmin}_{\bar{U}_{\infty}} \left(\frac{1}{N_{\mathcal{U}}} \sum_{i \in \mathcal{U}} \left(\bar{P}_i - \hat{P}_i(\phi, \bar{U}_{\infty}, \zeta, \bar{\gamma}_i) \right)^2 \right), \quad (4.2)$$

where \bar{P}_i and $\bar{\gamma}_i$ are the one-minute-averaged measured power capture and yaw angle of turbine i , respectively and \hat{P}_i is the power capture predicted by FLORIS.

3. The wake recovery factor is estimated using a five-minute-average of the turbine power measurements, as

$$\zeta = \operatorname{argmin}_{\bar{\zeta}} \left(\frac{1}{N_T} \sum_{i=1}^{N_T} \left(\bar{\bar{P}}_i - \hat{P}_i(\phi, U_{\infty}, \bar{\zeta}, \bar{\bar{\gamma}}_i) \right)^2 \right), \quad (4.3)$$

with N_T being the number of turbines, $\bar{\bar{P}}$ and $\bar{\bar{\gamma}}$ are vectors of length N_T containing the five-minute-averaged measured power signals and yaw angles, respectively and \hat{P} a vector of length N_T with the estimated power signal of each turbine according to FLORIS. Note that ζ is estimated using five-minute averages rather than one-minute averages to reduce variance and because ζ inherently varies much slower with time than the other variables. It is important to mention that ζ is only estimated when there is sufficient information in the measurements to do so. This relates to the *observability* \mathcal{O} of the situation. A more elaborate analysis concerning observability is performed in Chapter 3 of this dissertation.

The degree of observability \mathcal{O} for the various wind directions of the 6-turbine farm is shown in Figure 4.7. The top colormap shows to what extent we can reconstruct ζ from the measurements available in a steady-state situation. A value of $\mathcal{O} = 0$ implies that it can in no way be derived from the measurements, while a value of $\mathcal{O} = 1$ refers to the best-estimable situation.

The lower bar in Figure 4.7 saturates the observability to values of 0 and 1, with a threshold of 0.25 found empirically. In the farm control solution, if the observability over the past 400 seconds has been positive ($\mathcal{O} \geq 0.25$) for at least 80%

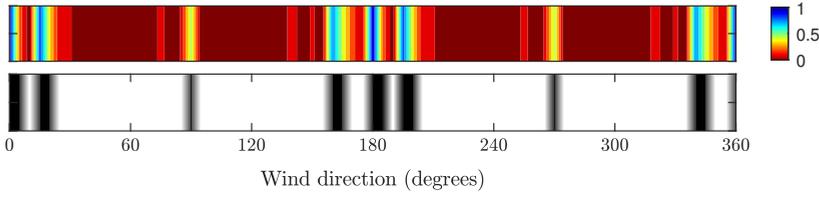


Figure 4.7: The observability of ζ when turbine power and wind direction measurements are available. The top bar shows the observability on a scale of 0 (not observable) to 1 (most observable). The bottom bar saturates to values of 0 (white) and 1 (black), with a threshold at 0.25. The observability is shown for the complete rose of wind directions. These plots are produced under a true ζ of 0.06 and $U_\infty = 8$ m/s. It is confirmed by simulation that the observability does not significantly vary over ζ and U_∞ .

of the time (black zones)¹ and the ambient conditions and control settings have not changed significantly in the last 400 seconds (i.e., a steady-state situation has arisen), then the estimate for ζ is updated. If not, then the wake recovery factor ζ is assumed to be equal to the last estimated value. It is confirmed by simulation that the observability does not significantly vary over ζ and U_∞ .

4.4.2. Real-time control setpoint optimization

After model adaptation, the turbine yaw angles are optimized in a robust manner following Rott et al. (2018) for maximum steady-state wind farm power production assuming a standard deviation on the wind direction of 2.5° in accordance to simulation data, as²

$$\gamma = \arg \max_{\tilde{\gamma}} \left(\mathbb{E} \left(\sum_{i=1}^{N_T} \hat{P}_i(\rho, U_\infty, \zeta, \tilde{\gamma}_i) \right) \right). \quad (4.4)$$

This optimization now contains \mathbb{E} denoting the expected value, since ρ is a Gaussian probability distribution of the wind direction with mean ϕ and a standard deviation of 2.5° . The optimal yaw setpoints are collected in a look-up table (LUT). Following this optimization, FLORIS assigns strong jumps in the yaw angle for small changes in ϕ , U_∞ and ζ as being optimal. These angles are therefore smoothed in post-processing using a 2D Gaussian distribution along ζ and ϕ with standard deviations of 0.04 and 3° . Note that the amount of smoothing necessary has a strong correlation with the variability of the ambient conditions. A more elaborate study would be necessary to determine the degree of smoothing that yields the best behavior.

The smoothed yaw setpoints for the 6-turbine case with a wind speed of $U_\infty = 8$ m/s and a low wake recovery factor of $\zeta = 0.07$ are shown in Figure 4.8. These setpoints are largely insensitive to the wind speed in region 2 operation (Kanev, 2020). The wind direction in Figure 4.8 is plotted along the x-axis, where 0° implies wind flowing from west to east and 90° implies wind flowing from south to north. In the 6-turbine layout,

¹The averaging time and the percentage thresholds here are found empirically.

²This cost function could straightforwardly be extended to include structural loads by, for example, penalizing the turbulence intensity in front of each turbine's rotor plane.

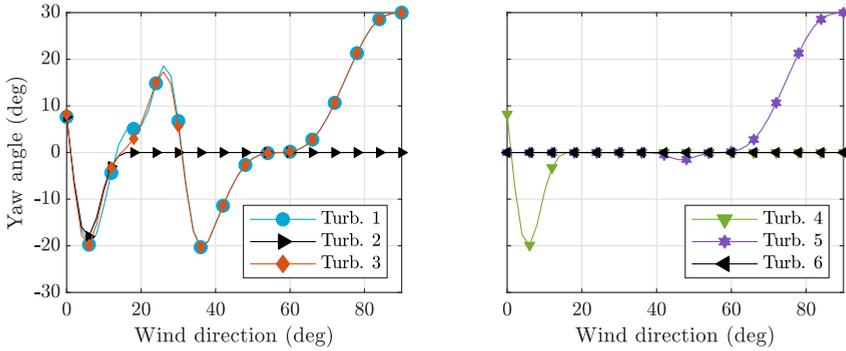


Figure 4.8: Optimized yaw angles as a function of wind direction for $\zeta = 0.07$.

this means that turbine 1 is always upstream and turbine 6 is always downstream. Hence, turbine 1 experiences a lot of yaw misalignment, while turbine 6 remains aligned over the entire wind range. It should be noted that, as ζ increases, wake losses diminish and the optimal yaw misalignment angles decrease. Essentially, there is less to be gained at downstream turbines.

4.4.3. An overview

The closed-loop wind farm control algorithm is synthesized by combining the estimator from Section 4.4.1 with the optimizer from Section 4.4.2. A pseudo-code is given in Algorithm 1, where the control setpoints are updated every 20 s.

4.5. Simulation results

In this section, the controller synthesized in Section 4.4 is tested in a high-fidelity simulation and compared to a locally greedy wind farm controller. In the latter, the yaw angles of the turbines are equal to the ambient wind direction dictated inside SOWFA. The closed-loop wind farm controller is subjected to a stress test in which the inflow varies strongly over time, being the SOWFA simulation of case D in Table 4.1. In Section 4.5.1, the estimation submodule is assessed. Then, in Section 4.5.2, the optimization submodule of the controller is evaluated by looking at the energy yield of the turbines. Finally, in Section 4.5.3, the yaw actuator duty cycle and the structural loads on the turbine blades are investigated.

4.5.1. Model adaptation performance

The first component of the closed-loop wind farm controller synthesized in Section 4.4 is the *model adaptation* block, as shown in Figure 4.1. In this simulation, the measurements fed to the wind farm controller are the instantaneous turbine power and wind direction measurements, of which the latter are artificially perturbed by Gaussian noise with a standard deviation of 2° to mimic measurement noise. The 6-turbine wind farm experiences a wind field of which the inflow direction and wind speed change often over

Algorithm 1 Pseudo-code of the closed-loop wind farm control algorithm

```

 $k \leftarrow 0$ 
while wind farm controller enabled do

  %% Gather time-averaged measurements
   $P, \gamma, \phi \leftarrow$  SCADA measurements

  if  $t = 20$  s, 40 s, 60 s, 80 s, ... then
     $k \leftarrow k + 1$ 

    %% Determine the freestream wind direction and wind speed
     $\phi^k \leftarrow \frac{1}{N_T} (\phi_1 + \phi_2 + \dots + \phi_{N_T})$  % ...with  $N_T$  the total number of turbines
     $\mathcal{U} \leftarrow$  upstream turbine indices
     $U_\infty^k \leftarrow$  Equation 4.2

    %% Determine the observability of  $\zeta$ 
     $\mathcal{O}^k \leftarrow$  observability( $\phi^k, U_\infty^k$ )

    if  $\mathcal{O}^k \geq 0.25$  then
       $\bar{\mathcal{O}}^k \leftarrow 1$ 
    else
       $\bar{\mathcal{O}}^k \leftarrow 0$ 
    end if

    %% Determine whether a steady-state situation has arisen
     $\sigma_\phi \leftarrow \text{std}(\phi^{k-20}, \phi^{k-19}, \dots, \phi^k)$ 
     $\sigma_\gamma \leftarrow \max_i (\text{std}(\gamma_i^{k-20}, \gamma_i^{k-19}, \dots, \gamma_i^k))$  % With  $i$  the turbine number
    SteadyState  $\leftarrow \text{bool}(\sigma_\phi \leq 1.0^\circ \ \& \ \sigma_\gamma \leq 1.0^\circ)$ 

    %% Estimate  $\zeta$  using time-averaged observability
    if SteadyState &  $\text{mean}(\bar{\mathcal{O}}^k, \bar{\mathcal{O}}^{k-1}, \dots, \bar{\mathcal{O}}^{k-20}) \geq 0.80$  then
       $\zeta \leftarrow$  Equation 4.3
    else
       $\zeta^k \leftarrow \zeta^{k-1}$ 
    end if

    %% Determine optimal yaw setpoints
     $\gamma \leftarrow \text{LookupYawTable}(\phi^k, U_\infty^k, \zeta^k)$ 

  end if
end while

```

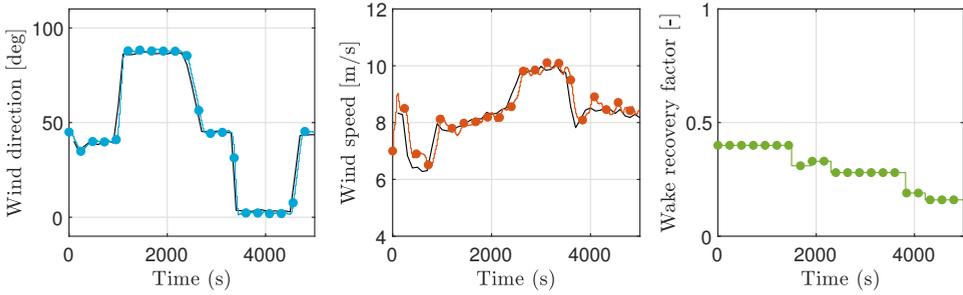


Figure 4.9: Overview of the estimator’s performance in reconstructing the ambient conditions. The true values are shown as solid black lines. The variables estimated by the controller are shown as solid colored lines.

time, as shown in Figure 4.9. In this figure, the solid black lines show the true values from SOWFA, while the solid colored lines show the estimated values according to the controller. The wake recovery factor ζ is initialized at a high value of 0.40 to enforce conservatism in the assigned control setpoints (small yaw misalignment angles) until there is sufficient information to derive a correct estimate for ζ from the measurements.

Taking a closer look at Figure 4.9, it becomes clear that both the wind direction and the wind speed are estimated accurately and consistently. Note that wind direction is by far the most important variable to estimate, as the optimal yaw setpoints are most sensitive to this variable and less sensitive to the wind speed and wake recovery factor (Kanev, 2020). Note that the estimate for the wake recovery factor ζ is only updated when the situation is sufficiently *observable*, as defined in Section 4.4.1. The fluctuations in wind direction often lead to situations with little to no wake interaction, yielding a low observability and hence ζ is constant until about 1500 s into the simulation. Then, it is consistently estimated to be around 0.17 – 0.31. In FLORIS, a high wake recovery factor leads it to predict high wake recovery, which then leads to small yaw misalignment setpoints. A low wake recovery factor leads to large wake losses and therefore higher yaw misalignment setpoints.

The five-minute-averaged flow field and turbine power signals for SOWFA and FLORIS are shown in Figure 4.10. One can see that the model shows a mediocre match in the flow fields at $t = 900$ s. Namely, the freestream wind speed and wind direction are in the right ballpark, but the prediction lacks in the far wake and in the turbine power signals. This is not in the least due to the conservative initial value for ζ . The model prediction improves at 1800 s and onward due to the adaptation of ζ . Furthermore, most of the model errors seem to originate behind the second row of turbines. One possible explanation for this is the lack of secondary wake steering effects in FLORIS, as discussed in Martínez-Tossas et al. (2019). Moreover, as the turbine power signals are used for model adaptation, it is no surprise that the turbine power signals match very well between SOWFA and FLORIS. An interesting difference between SOWFA and FLORIS is that FLORIS predicts many situations of symmetry, in which turbines are predicted to capture an equal amount of power. In SOWFA, however, the turbulent inflow unavoidably gives rise to differences in

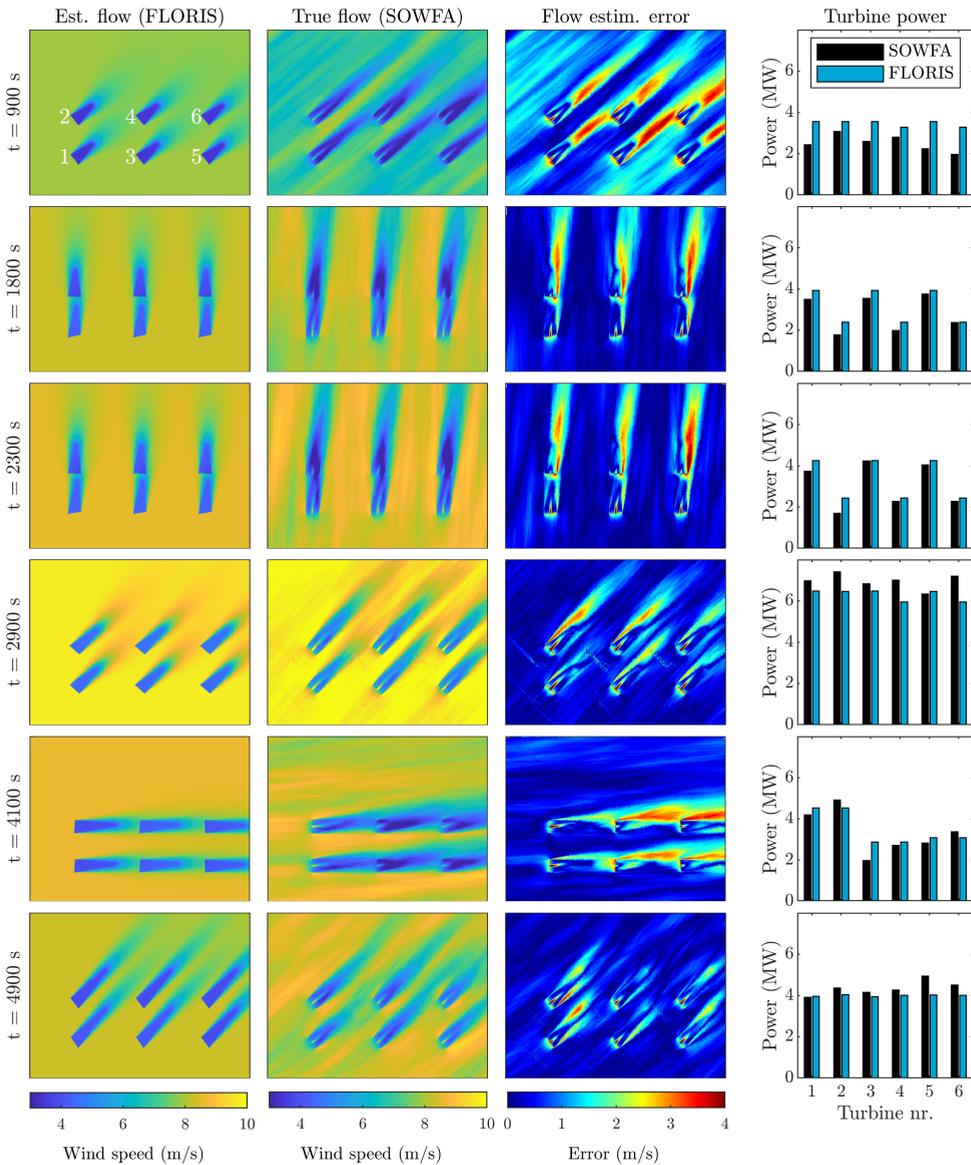


Figure 4.10: The estimation performance of the closed-loop wind farm controller. The five-minute-averaged true farm's flow field and power signals from SOWFA are compared to those estimated by FLORIS. This figure clearly shows that the wind farm controller accurately predicts the power signals and flow fields of SOWFA, even though there are significant discrepancies in the model's underlying equations.

power capture between turbines.

Note that ζ deviates significantly from the physical turbulence intensity, as it is com-

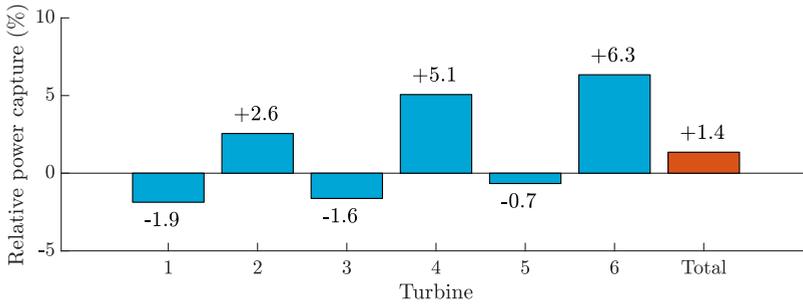


Figure 4.11: Change in turbine electrical power capture compared to the case in which turbines are controlled in the traditional, locally greedy manner. The values here are averaged over the 5000 s of simulation. The turbines that are most often upstream (1, 3 and 5) experience most losses, as they are most often misaligned with the inflow. The largest gains are attained by the turbines that operate most often in a waked inflow (2, 4 and 6). On average, wake steering led to an increase in the total wind farm energy yield of 1.4% in the 5000 s of simulation.

monly defined as. This allows a further reduction of the difference between SOWFA and FLORIS, as successfully demonstrated in the simulation study at hand. Despite the discrepancies in the fundamental equations underlying FLORIS and SOWFA, the model adaptation algorithm provides accurate and consistent estimates of the inflow conditions and the wake recovery factor.

4.5.2. Setpoint optimization performance

With the ambient conditions estimated, the second component of the closed-loop wind farm controller is setpoint optimization (recall Figure 4.1). In this chapter, that consists of the optimization of the turbine yaw misalignment angles to maximize the power extraction of the wind farm.

In Figure 4.11, the relative power capture of each turbine normalized to the locally greedy wind farm controller scenario is shown, averaged over the 5000 s of simulation. Since the wind changes from an inflow from *west to east* to an inflow from *south to north* and anywhere in between, turbines 1, 2, 3 and 5 are most often upstream, while turbines 4 and 6 are mostly downstream. This explains the energy loss in turbines 1, 3 and 5 and it also explains the energy gain in turbines 4 and 6. Over the total 5000 s of simulation, the energy yield is 1.4% higher with the closed-loop wind farm controller compared to the baseline case.

Note that an improvement in energy yield of 1.4% is lower than most values cited in the literature (Boersma et al., 2017), as those studies typically only focus exclusively on situations with significant wake losses. In this simulation, at several time instants, there is little to no wake interaction. This is demonstrated in Figure 4.12, showing the relative gain in wind-farm-wide energy yield over time. In this figure, it is seen that the increase in instantaneous wind-farm-wide power yield varies between -4% and $+11\%$, depending on the inflow conditions. For the wake-loss-heavy scenario in the time window of 1800 to 2300 s, a total increase in energy yield of 7.3% is noted. Similarly, for the second

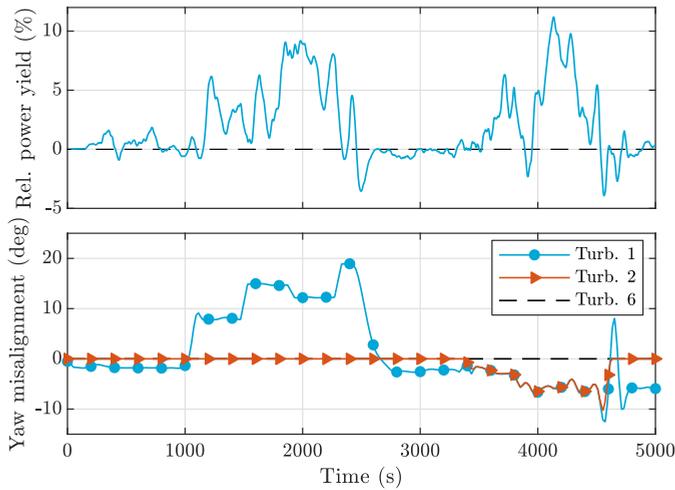


Figure 4.12: The top plot shows the wind farm power capture over time. While the time-averaged change in power production is 1.4% (Figure 4.11), the instantaneous change in power capture varies between -4% and $+11\%$ throughout the simulation. Furthermore, the yaw angles of a subset of turbines over time are shown in the bottom plot. Turbine 1 experiences misalignments of -11° to $+20^\circ$ as it is the most upstream turbine, whereas turbine 6 maintains zero misalignment, always being the most downstream machine throughout the simulation.

wake-loss-heavy scenario in the time window of 4000 to 4500 s, an increase in energy of 6.0% is noted. Moreover, energy losses appear for short periods of time throughout the simulation, mostly prevailing an increase in energy yield due to yaw steering. This is because the benefit of misaligning upstream turbines is not noticed until the flow has propagated to the downstream turbines, which takes approximately 100 s in these situations. The corresponding yaw angles for turbines 1, 2 and 6 are also displayed in the figure, showing misalignment angles of up to 20 degrees. As the wake of turbine 6 never impinges another turbine, its yaw angle remains zero throughout the simulation. In total, a gain in power production of 1.4% over the 5000 s of simulation despite the large discrepancies between FLORIS and SOWFA is still a very promising (and more realistic) estimate of the true potential of wake steering.

A final remark is that FLORIS somewhat underpredicts the amount of wake displacement achieved due to a yaw misalignment, as seen in Figure 4.10. Therefore, the proposed closed-loop wind farm control solution is somewhat conservative, assigning relatively small yaw angle setpoints to the turbines. More energy than presented currently may be harvested by refining the FLORIS model for wake steering.

4.5.3. A deeper look into the yaw actuator duty cycle and structural loads

Wake steering shows to be very promising in increasing the energy yield of a wind farm. In contrast to the energy gains, the effects of wake steering on the actuator duty cycle and the structural loads on the turbines remain unclear. This subsection addresses these two

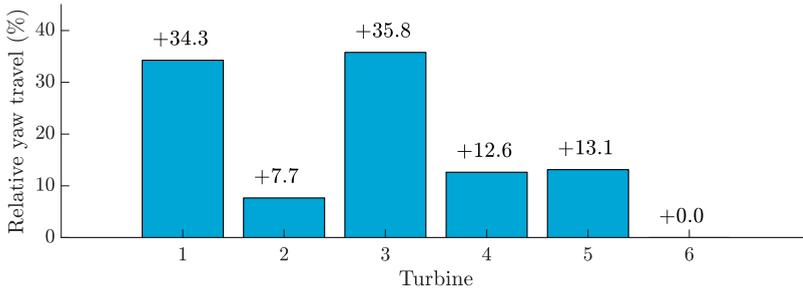


Figure 4.13: Total distance of yaw travel throughout the 5000 s simulation compared to the locally greedy case. This figure shows the additional load put on the yaw actuator to track the assigned yaw setpoints. Turbines 1-5 are all misaligned at some point in time, thereby increasing the yaw travel compared to locally greedy control. Since turbine 6 is never misaligned, it has the same yaw travel as the baseline simulation.

topics.

Yaw actuator duty cycle under wake steering

The change in the yaw actuator duty cycle (*yaw travel*) of each turbine is shown in Figure 4.13. From this figure, it is clear to see that wake steering has a noticeable influence on the yaw actuator duty cycle in this simulation. Specifically, for the upstream turbines, an increase in yaw travel of up to 36% is seen. Furthermore, even for the second row of turbines, an increase in yaw travel of 8 – 13% is seen. Also, turbine 6 has no additional yaw travel as it always operates most downstream of the wind farm for the simulated wind conditions. The amount of yaw travel relates back to Figure 4.8, in which relatively large gradients can be seen for small changes in the wind direction. One may reduce the yaw travel by further smoothing the optimized yaw angles from Figure 4.8, but this may go at the loss of energy yield. In a practical controller implementation, a trade-off must be made according to the yaw actuator limits, the wind farm layout and the wind rose of the wind farm.

Fatigue loads on the blade roots

The bending moments around the blade root are calculated for each blade. The damage-equivalent loads (DELs) of the out-of-plane bending moments at the blade roots are then calculated following the Palmgren Miner's rule (Cosack, 2010), as

$$\text{DEL} = \sqrt[m]{\frac{\sum_{i=1}^n \Delta S_i^m \cdot N_i}{N_{\text{ref}}}}. \quad (4.5)$$

In this equation, m is the inverse of the material Wöhler slope and N_{ref} is a reference number for the total amount of cycles taken to be 1 here, S_i is the mean load range value for a particular bin and N_i is the number of occurrences within the bin. The load cycles are calculated following the popular rainflow counting method (ASTM International, 2017). Note that the blades of the DTU 10MW turbine are manufactured with a

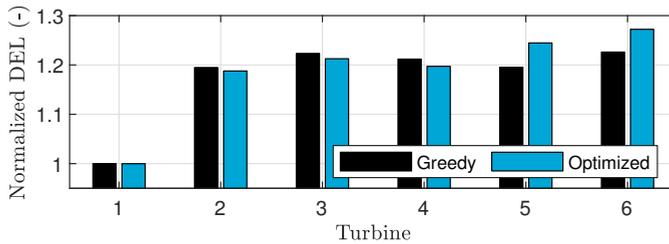


Figure 4.14: Damage equivalent loads for the blade root out-of-plane bending moments, normalized with respect turbine 1 in greedy operation.

mix of glass fiber, carbon fiber and balsa (Bak et al., 2012). Hence, the DEL values are evaluated with both $m = 10$ (glass fiber) and with $m = 14$ (carbon fiber) and the highest value of the two will be shown.

The DELs normalized by the values of turbine 1 under greedy control are plotted in Figure 4.14. The loads between the three blades are very similar and hence the blade-averaged DELs are shown. From this figure, it is seen that the DELs mostly decrease for turbines 1-3 with the closed-loop wind farm controller compared to baseline operation and increase for turbines 4-6. For turbines 1-3, the decrease in DELs is expected to be due to the reduction in effective wind speed due to the applied yaw misalignments. Turbines 4-6 operate in waked flow more often and the increase in DELs is expected to be due to the increase in the rotor-effective turbulence intensity and wind speed.

However, generally, wake steering seems to have a relatively small effect on the blade root out-of-plane bending moments in this simulation study. This may be explained by literature that suggests that the blade loads can both increase and reduce as a result of yaw misalignment, depending on the yaw direction and the amplitude (Reyes et al., 2019). Moreover, this publication also indicates that there is a delicate balance between the change in loads due to yawing a turbine and due to the change in the wake profile as a result of the yawing of an upstream turbine. Though, simulations with a higher-fidelity aero-elastic models and physical experiments are necessary to further solidify such statements.

4.6. Conclusions

In this chapter, a novel, closed-loop wind farm controller was proposed. This control solution relies on the popular, steady-state, computationally efficient FLORIS surrogate wind farm model of the wind farm. The controller consists of two parts. Firstly, FLORIS is used to estimate the freestream wind direction, wind speed and the wake recovery factor, supported by a theoretical measure of observability to prevent the estimation of quantities about which no information is available. Secondly, FLORIS is leveraged to optimize the turbine yaw setpoints for energy yield maximization.

This closed-loop and model-based wind farm control solution was tested in a high-fidelity simulation subjected to a time-varying inflow, being the first of its kind in the literature. The wind direction and wind speed in the simulation contain strong changes

to stress-test the controller. Compared to baseline operation, a total time-averaged gain in energy yield of 1.4% was found for a virtual 6-turbine offshore wind farm with 10MW turbines. Moreover, for particular time windows, gains in the energy yield of up to 11% were noted, agreeing well with studies in the literature (Boersma et al., 2017). Furthermore, from an actuator duty cycle perspective, the yaw travel of the turbines increased up to 36%. Additionally, the damage-equivalent loads of the blade root out-of-plane bending moments did not change significantly compared to the baseline controller.

The results presented in this chapter highlight the potential of the proposed controller, even when subjected to time-varying inflow conditions, addressing an important phenomenon in real wind farms. This solidifies the proposed control solution as the first realistic, closed-loop wind farm control solution for yaw-based wake steering.

A number of recommendations can be made for future research. Firstly, the proposed control solution was stress-tested in this chapter. To get a realistic idea of the effect of wake steering on the annual energy production, one would have to perform experiments with realistic wind profiles, e.g., generated from real measurement data. This is forthcoming in van der Hoek et al. (2020). Moreover, the simulations would benefit from a larger simulation domain, even though this increases computational cost. Secondly, the optimal update frequency of the yaw setpoints was not considered in this chapter and, at large, has been addressed insufficiently in the literature (Kanev, 2020). Thirdly, continued work on surrogate wind farm modeling (e.g., Martínez-Tossas et al., 2019; Soleimanzadeh et al., 2014) should further improve wind farm controller performance, going hand-in-hand with developments in estimation (e.g., Chapter 5 of this dissertation) and optimization algorithms (e.g., Annoni et al., 2018a). Important topics for surrogate modeling include wake propagation, time-varying inflow, spatially varying inflow, atmospheric stability effects and local variations in wind characteristics due to, e.g., terrain effects. Finally, while considered a high-fidelity testing environment, SOWFA remains a simulation model and field experiments are essential to further increase confidence in the proposed algorithm and, at large, wake steering for power maximization in wind farms.

5

ONLINE MODEL CALIBRATION FOR A DYNAMIC SURROGATE MODEL

ABSTRACT

Wind farm control often relies on computationally inexpensive surrogate models to predict the dynamics inside a farm. However, the reliability of these models over the spectrum of wind farm operation remains questionable due to the many uncertainties in the atmospheric conditions and tough-to-model dynamics at a range of spatial and temporal scales relevant for control. A closed-loop wind farm control framework is proposed in which a simplified wind farm model is calibrated and used for optimization in real time. This chapter presents a joint state-parameter estimation solution with an Ensemble Kalman filter at its core, which calibrates the surrogate wind farm model to the actual atmospheric conditions. The estimator is tested in high-fidelity simulations of a nine-turbine wind farm. Using exclusively SCADA measurements, the adaptability to modeling errors and mismatches in atmospheric conditions is shown. Convergence is reached within 400 seconds of operation, after which the estimation error in flow fields is negligible. At a low computational cost of 1.2 s on an 8-core CPU, this algorithm shows comparable accuracy to the state of the art from the literature while being approximately two orders of magnitude faster.

5.1. Introduction

Over the past decades, global awakening on climate change and the environmental, political and financial issues concerning fossil fuels have been catalysts for the growth of the renewable energy industry. As the primary energy demand in Europe is projected to decrease by 200 million tonnes of oil equivalent from 2016 to 2040, there is an additional shift in the energy source used to meet this demand (International Energy Agency, 2017). Shortly after 2030, onshore and offshore wind energy are projected to become the main source of electricity for the European Union. By then, about 80% of all new capacity added is projected to come from renewable energy sources, enabled by a favorable political climate.

While these developments have clear benefits, an important problem with wind energy is that the rotational speed of most commercial turbines is decoupled from the electricity grid frequency via each turbine's power electronics (Aho et al., 2012). As the current grid-connected fossil fuel plants are replaced by non-synchronous renewable energy plants, the inertia of the electricity grid will decrease, making it less stable and more prone to machine damage and blackouts (Ela et al., 2014). Therefore, there is a strong need for wind farms and other renewables to provide ancillary grid services. Wind farm control aimed at increasing the grid stability is more commonly defined as active power control (APC). In APC, the power production of a wind farm is regulated to meet the power demand of the electricity grid, which may change from second to second.

Existing literature on wind farm control has mainly focused on maximizing the power capture (e.g., Gebraad and van Wingerden, 2015; Gebraad et al., 2016; Munters and Meyers, 2017; Rotea, 2014). Though, literature on APC has been receiving an increasing amount of attention (e.g., Boersma et al., 2017; Fleming et al., 2016; van Wingerden et al., 2017). The main challenges in wind farm control are the large time delays caused by the formation of wakes, the many uncertainties in the atmospheric conditions and the questionable reliability of surrogate wind farm models over the wide spectrum of wind farm operation. See Boersma et al. (2017) and Knudsen et al. (2015) for state-of-the-art overviews of control and control-oriented modeling for wind farms. While there has been success with model-free methods for power maximization (e.g., Rotea, 2014), it is unclear to what degree such methods can be used for power forecasting. Furthermore, model-free methods typically have long settling times, making them intractable for APC. On the other hand, for model-based approaches, the aforementioned challenges make it impossible for any model to reliably provide power predictions in an open-loop setting. Hence, a model-based approach in which a surrogate wind farm model is actively adjusted to the present conditions is a necessity for reliable and computationally tractable APC algorithms. This closed-loop wind farm control framework, consisting of three components, is shown in Figure 5.1.

The first component of the closed-loop framework is a computationally inexpensive surrogate model that accurately predicts the power production of the wind farm ahead in time, on a time-scale relevant for control. The most commonly used surrogate models in wind farm control are steady-state models, which are heuristic and neglect all temporal dynamics (Boersma et al., 2017). While some of these models have shown success in

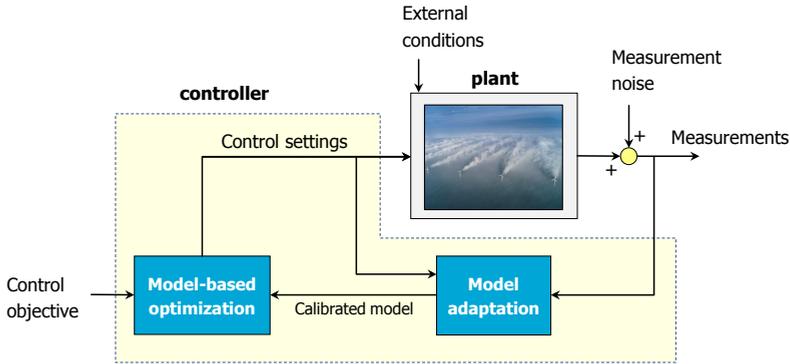


Figure 5.1: Closed-loop wind farm control framework. Measurements (e.g., SCADA or Light Detection And Ranging (LiDAR) data) are fed into the controller. First, the state of the surrogate wind farm model is estimated to represent the actual atmospheric and turbine conditions inside the wind farm. Secondly, using the calibrated model, an optimization algorithm determines the control policy (e.g., yaw angles) for all turbines. This control policy may be a set of constant operating points, but can also be time-varying, depending on whether the surrogate model is time-varying and the employed optimization algorithm. The photograph of the wind farm is from Christian Steiness.

wind tunnel tests (e.g., Schreiber et al., 2017) and field tests (e.g., Fleming et al., 2017a,b) for power maximization, the actuation frequency is limited to the minutes-scale, since the flow and turbine dynamics are predicted on the minute-scale. Furthermore, time-ahead predictions with these models are limited to the steady state, limiting their use for APC. There is a smaller yet significant number of dynamic surrogate wind farm models (e.g., Boersma et al., 2018a; Munters and Meyers, 2017; Shapiro et al., 2017a), which attempt to include the dominant temporal dynamics inside the farm. These models can be used for control on the seconds-scale and furthermore allow time-ahead predictions, some even under changing atmospheric conditions. Specifically, the dynamic surrogate model employed in Shapiro et al. (2017a) is computationally feasible, but only models the flow in one dimension and furthermore allows no turbine yaw or changes in the wind direction, limiting its applicability. Furthermore, the dynamical model in Munters and Meyers (2017) has shown success for closed-loop wind farm control applications, but it is too computationally costly for any kind of real-time control and the authors present their results solely as a benchmark case. In the work presented here, the model described in Boersma et al. (2018a) is used, which is a two-dimensional LES code with wind farm control as its main objective. This dynamic surrogate model, named “Wind-FarmSimulator” (WFSim), includes yaw and axial induction actuation, turbine-induced turbulence effects and spatially and temporally varying inflow profiles, with a moderate computational cost.

The second component of the closed-loop wind farm framework is an algorithm that adjusts the surrogate model’s parameters to improve its accuracy online using flow and/or turbine measurements (e.g., SCADA data, LiDAR measurements, met masts). In terms of control, this turns into a joint estimation problem, in which both the model

state and a subset of model parameters are estimated online. Currently, the optimization algorithms presented in Munters and Meyers (2017) and Vali et al. (2017) have assumed full state knowledge, conveniently ignoring the step of model adaptation. Literature on state reconstruction and model calibration for dynamical wind farm models is sparse, limited to linear low-order models and/or common estimation algorithms. Gebraad et al. (2015) designed a traditional Kalman filter (KF) for their low-fidelity model, showing marginal improvements compared to optimization using a static model. Shapiro et al. (2017a) present a one-dimensional dynamic wake model used with receding horizon control for secondary frequency regulation, using an estimation algorithm following Doekemeijer et al. (2016). Furthermore, Iungo et al. (2015) used dynamic mode decomposition to obtain a reduced-order model of the wind farm dynamics, which was then combined with a traditional KF for state estimation. To the best of the authors' knowledge, none of these methods have explored more sophisticated models such as WFSim and often only use simple state estimation algorithms that are lacking in terms of accuracy and computational tractability.

The third component of the closed-loop framework is an optimization algorithm, which typically is a gradient-based or nonlinear optimization algorithm (e.g., Gebraad et al., 2016) for steady-state models and a predictive optimization method for dynamical models (e.g., Goit and Meyers, 2015; Siniscalchi-Minna et al., 2018; Vali et al., 2017). A more in-depth discussion on optimization algorithms is out of the scope of this chapter.

The focus of this chapter is on a model adaptation algorithm for WFSim, which trades off estimation accuracy with computational complexity. In previous work (Doekemeijer et al., 2016, 2017), state estimation using flow measurements downstream of each turbine has shown success using an Ensemble KF (EnKF), with a computational cost several orders of magnitude lower than traditional KF methods. The main contributions of this chapter relative to Doekemeijer et al. (2016, 2017) are:

- the additional adaptation to a mismatch in atmospheric conditions (specifically, the ambient wind speed and turbulence),
- the option to use turbine power signals in addition to, or instead of, flow measurements,
- a further reduction in the computational complexity,
- a comparison of the EnKF with the state of the art in the literature.

The structure of this chapter is as follows. In Section 5.2, the surrogate wind farm model WFSim is introduced. In Section 5.3, a time-efficient, online model calibration algorithm for dynamical wind farm models is detailed. This calibration algorithm is validated and compared with standard algorithms in the literature in high-fidelity simulations in Section 5.4. The chapter is concluded in Section 5.5.

5.2. The surrogate model

The framework of Figure 5.1 requires a surrogate model of the wind farm. In this chapter, that is the WindFarmSimulator (WFSim) model presented by Boersma et al. (2018a). This model is particularly suited as it includes both yaw and axial induction actuation and

yields a relatively high accuracy with a relatively low computational cost.¹ The scope of this section is to give a summary of the surrogate model, rather than a full derivation and motivation of the assumptions made. The reader is referred to Boersma et al. (2018a) for more information.

Fundamentally, WFSim is based on the two-dimensional unsteady incompressible Navier-Stokes (NS) equations. The surrogate model can be described completely by the flow and rotor dynamics in a horizontal plane at hub height. WFSim deviates from a traditional 2D NS model in two ways. Firstly, the diffusion term is neglected, as it plays a negligible role due to the low viscosity of air. Secondly, the dissipation term in the lateral direction in the continuity equation is multiplied by a factor 2 to approximate flow dissipating in the vertical flow dimension. Other vertical flow contributions such as vertical meandering and shear are neglected. The subgrid-scale model is formulated using an eddy-viscosity assumption in combination with Prandtl's mixing length model. The mixing length is parametrized as a function of the spatial location, increasing linearly with distance from the downstream rotor, starting at zero at distance d' downstream and peaking at distance d , where ℓ_s defines the slope of the mixing length. Basically, the larger ℓ_s , the quicker wakes recover to their freestream properties. Furthermore, the turbines are modeled using the non-rotating (static) actuator disk model (ADM), projected onto the 2D plane at hub height. The turbine is assumed to be a rigid object applying a two-dimensional force vector on the flow. Both the turbine forcing term and the turbine power output are scaled by tuning factors c_f and c_p , respectively, to account for unmodeled effects. Together with the three parameters from the turbulence model, this leads to a total of five tuning parameters.

These NS equations are solved over a spatially and temporally discretized domain (Boersma et al., 2018a). Dirichlet boundary conditions for the longitudinal and lateral velocity are applied on one side of the grid for inflow, while Neumann boundary conditions are applied on the remaining sides for the outflow. The surrogate model reduces to a nonlinear discrete-time deterministic state-space model, as

$$\begin{aligned}x_{k+1} &= f(x_k, q_k), \\z_k &= h(x_k, q_k),\end{aligned}$$

where $x_k \in \mathbb{R}^N$ is the system state at time k , which is a column vector containing the collocated longitudinal flow velocity at each cell in the domain $\tilde{u}_k \in \mathbb{R}^{N_u}$, the lateral flow velocity at each cell in the domain $\tilde{v}_k \in \mathbb{R}^{N_v}$ and the pressure term at each cell in the domain $\tilde{p}_k \in \mathbb{R}^{N_p}$, with $N = N_u + N_v + N_p$ and $N_u \approx N_v \approx N_p \approx \frac{1}{3}N$. The state x_k is formulated as

$$x_k^T = [\tilde{u}_k^T \quad \tilde{v}_k^T \quad \tilde{p}_k^T].$$

Empirically, good results have been achieved with cell dimensions of about 30 – 50 m in width and length, resulting in N with a typical value on the order of $10^3 - 10^4$ for six-

¹Note that it is still uncertain what accuracy is necessary and what computational cost can be permitted for real-time closed-loop wind farm control.

to nine-turbine wind farms (e.g., Boersma et al., 2018a; Doekemeijer et al., 2016, 2017; Vali et al., 2017). Such a number of states may seem very small for LES simulations, yet is very high for control purposes. Furthermore, $q_k \in \mathbb{R}^O$ includes the system inputs, i.e., the turbine control settings: the turbine yaw angles γ_i and the thrust coefficients C'_{T_i} for $i = 1, \dots, N_T$, with N_T the number of turbines. The system outputs $z_k \in \mathbb{R}^M$ are defined by sensors. It can include, among others, flow field measurements ($z_k \subset x_k$) and power measurements. We define the integer $M_{u,v} \in \mathbb{Z}$ with $0 \leq M_{u,v} \leq M$ as the total number of flow field measurements. The nonlinear functions f and h are the state forward propagation and output equation, respectively.

The computational cost may vary from 0.02 s for two-turbine wind farms with $N = 3 \cdot 10^3$ states (e.g., in Doekemeijer et al., 2017), to 1.2 s for $N = 1 \cdot 10^5$ states for medium-sized wind farms (e.g., in Boersma et al., 2018a), for a single time-step forward simulation on a single desktop CPU core. The computational complexity of the model is what motivates the use of time-efficient estimation algorithms in this chapter and time-efficient predictive control methods for optimization in related work (Vali et al., 2017). Here, the limits of computational cost are explored to maximize model accuracy while still allowing real-time control. Note that research on the computational feasibility of optimization algorithms using WFSim is ongoing.

5.3. Online model calibration

Due to the limited accuracy of surrogate wind farm models and due to the many uncertainties in the environment, surrogate models often yield predictions with significant uncertainty of the wind flow and power capture inside a wind farm. Since control algorithms largely rely on such predictions, this may suppress gains or even lead to losses inside a wind farm. Unfortunately, higher-fidelity models are computationally prohibitively expensive for control applications. Hence, rather, lower-fidelity surrogate wind farm models are calibrated online using readily available measurement equipment.

In this section, first the challenges for real-time model calibration for the surrogate “WFSim” model described in Section 5.2 will be highlighted in Section 5.3.1. Secondly, a mathematical framework for recursive model state estimation will be presented in Section 5.3.2. Thirdly, a number of nonlinear state estimation algorithms are presented in Sections 5.3.3 to 5.3.5, building up from the industry standard to the state of the art in the literature. Finally, a robust, computationally efficient model calibration solution is synthesized in Section 5.3.6, which allows the simultaneous estimation of the boundary conditions, model parameters and the model states of WFSim in real time using readily available measurements from the wind farm.

Note that we will henceforth refer to the estimation of x as *state(-only) estimation*. The estimation of both model states and model parameters such as ℓ_s is referred to as *(joint) state-parameter estimation*.

5.3.1. Challenges

Online model calibration for WFSim is challenging for a number of reasons. First of all, the model is nonlinear and thus the common linear estimation algorithms cannot be used without linearization, which limits accuracy (Boersma et al., 2018a). Secondly, an estimation solution relying on WFSim is sensitive to instability when the model state sufficiently deviates from the continuity equation. Finally, the surrogate model typically has on the order of $N \sim 10^3 - 10^4$ states, which is extraordinarily high for control applications. Though, real-time estimation is a necessity for real-time model-based control and thus one needs to find a trade-off between accuracy while guaranteeing state updates at a low computational cost.

5.3.2. General formulation

This section summarizes the basics of the Kalman Filter (KF), which is the literature standard for state estimation in control. The goal of a KF is to recursively estimate the unmeasured states of a dynamical system through noisy measurements. Assumed here is a system (the wind farm) represented mathematically by a discrete-time stochastic state-space model with additive noise,

$$x_{k+1} = f(x_k, q_k) + w_k, \quad (5.1)$$

$$z_k = h(x_k, q_k) + v_k, \quad (5.2)$$

where k is the time index, $x \in \mathbb{R}^N$ is the unobserved system state, $z \in \mathbb{R}^M$ are the measured outputs of the system, $q \in \mathbb{R}^O$ and $w \in \mathbb{R}^N$ are the controllable inputs and process noise respectively that drive the system dynamics and $v \in \mathbb{R}^M$ is measurement noise. Furthermore, we assume w and v to be zero-mean white Gaussian noise with covariance matrices

$$\mathbb{E} \left[\begin{bmatrix} v_k \\ w_k \end{bmatrix} \begin{bmatrix} v_\ell^T & w_\ell^T \end{bmatrix} \right] = \begin{bmatrix} R_k & S_k^T \\ S_k & Q_k \end{bmatrix} \Delta_{k-\ell}, \quad \text{where } \Delta_{k-\ell} = \begin{cases} 1, & \text{if } k = \ell, \\ 0, & \text{otherwise,} \end{cases} \quad (5.3)$$

with \mathbb{E} the expectation operator. Estimates of the state x_k , denoted by $\hat{x}_{k|k}$, are computed based on measurements from the real system. Here, $\hat{x}_{k|\ell}$ means an estimate of the state vector x at time k , using all past measurements and inputs \mathcal{Z}_ℓ , as

$$\hat{x}_{k|\ell} = \mathbb{E}[x_k | \mathcal{Z}_\ell], \quad \text{with } \mathcal{Z}_\ell = z_0, z_1, z_2 \dots z_\ell, \quad q_0, q_1, q_2 \dots q_\ell. \quad (5.4)$$

State estimates are based on the internal model dynamics and the measurements, weighted according to their probability distributions. We aim to find an optimal state estimate, in which optimality is defined as unbiasedness, $\mathbb{E}[x_k - \hat{x}_k] = 0$ and when the variance of any linear combination of state estimation errors (e.g., the trace of $\mathbb{E}[(x_k - \hat{x}_k)(x_k - \hat{x}_k)^T]$) is minimized.

In reality, the assumed model described by f and h always has mismatches with the true system and the assumptions in Equation 5.3 often do not hold. Further, the matrices Q_k , R_k and S_k are usually not known and rather considered tuning parameters, used to shift the confidence levels between the internal model and the measured values. For

$R \ll Q$, estimations will heavily rely on the measurements, while for $Q \ll R$, estimations will mostly rely on the internal model. Kalman filtering remains one of the most common methods of recursive state estimation. KF algorithms typically consist of two steps, namely:

1. A state and output forecast, including their uncertainties (covariances):

$$\hat{x}_{k|k-1} = E[f(x_{k-1}, q_{k-1}) + w_{k-1} | \mathcal{Z}_{k-1}], \quad (5.5)$$

$$\hat{z}_{k|k-1} = E[h(x_k, q_k) + v_k | \mathcal{Z}_{k-1}], \quad (5.6)$$

$$P_{k|k-1}^x = \text{Cov}(x_k, x_k | \mathcal{Z}_{k-1}) = E[(x_k - \hat{x}_{k|k-1})(x_k - \hat{x}_{k|k-1})^T], \quad (5.7)$$

$$P_{k|k-1}^z = \text{Cov}(z_k, z_k | \mathcal{Z}_{k-1}) = E[(z_k - \hat{z}_{k|k-1})(z_k - \hat{z}_{k|k-1})^T], \quad (5.8)$$

$$P_{k|k-1}^{xz} = \text{Cov}(x_k, z_k | \mathcal{Z}_{k-1}) = E[(x_k - \hat{x}_{k|k-1})(z_k - \hat{z}_{k|k-1})^T]. \quad (5.9)$$

In Equations 5.5 and 5.6, $\hat{x}_{k|k-1}$ and $\hat{z}_{k|k-1}$ are the forecasted system state vector and measurement vector, respectively.

2. An analysis update of the state vector, where the measurements are fused with the internal model:

$$L_k = P_{k|k-1}^{xz} \cdot (P_{k|k-1}^z)^{-1} \quad (5.10)$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + L_k (z_k - \hat{z}_{k|k-1}), \quad (5.11)$$

$$P_{k|k}^x = \text{Cov}(x_k, x_k | \mathcal{Z}_k) = P_{k|k-1}^x - L_k P_{k|k-1}^z L_k^T. \quad (5.12)$$

Here, $(P_{k|k-1}^z)^{-1}$ in Equation 5.10 is the pseudo-inverse of $P_{k|k-1}^z$, since this matrix is not necessarily invertible.

Traditionally, state estimation for linear dynamic models is done using the linear KF (Kalman, 1960). However, this is not a viable option here, as the surrogate wind farm model is nonlinear. Rather, a number of nonlinear KF variants are looked upon.

5.3.3. Extended Kalman filter (ExKF)

Linearization of the surrogate wind farm model is the most popular and straight-forward solution to the issue of model nonlinearity, as done in the Extended KF (ExKF). The ExKF has shown success in academia and industry (Wan and Merwe, 2000) and is perhaps the most popular nonlinear KF. However, it has a number of disadvantages. As described in Section 5.3.1, model linearization is troublesome. Furthermore, for surrogate models with many states such as WFSim, the ExKF has an additional challenge: computational complexity. The operation in Equation 5.10 includes a matrix inversion with a computational complexity of $\mathcal{O}(M^3)$ and the ExKF furthermore includes two matrix multiplications each with a complexity of $\mathcal{O}(N^3)$. As there are significantly fewer measurements than states ($M \ll N$) for the problem at hand, these matrix multiplications dominate the computational cost. The ExKF has a CPU time in the order of 10^1 s for a two-turbine wind farm, which may be too large for our purposes. To reduce computational cost in the

ExKF, the surrogate model and/or the covariance matrix P have to be simplified. This is not further explored here. Instead, two KF approaches will be explored that use the nonlinear system directly for forecasting and analysis updates. Doing so, we circumvent the problems with linearization and additionally better maintain the true covariance of the system state.

5.3.4. The Unscented Kalman filter (UKF)

The Unscented Kalman filter (UKF) relies on the so-called “unscented transformation” to estimate the means and covariance matrices described by Equations 5.5 to 5.9. The conditional state probability distribution of x_k knowing \mathcal{Z}_k is again assumed to be Gaussian. In the UKF, firstly a number of sigma points (also referred to as “particles”) are generated such that their mean is equal to $\hat{x}_{k|k}$ and their covariance is equal to $\text{Cov}(x_k, x_k)$. Secondly, each particle is propagated through the nonlinear system dynamics (f, h). Thirdly, the mean and covariance of the forecasted state probability distribution is again approximated by a weighted mean of these forecasted sigma points (Wan and Merwe, 2000).

Mathematically, we define the i^{th} particle as $\boldsymbol{\psi}_{k|\ell}^i \in \mathbb{R}^N$, which is a realization of the conditional probability distribution of x_k given \mathcal{Z}_ℓ . The UKF follows a very similar forecast and analysis update approach as the traditional KF in Equations 5.5 to 5.12, yet applied to a finite set of particles (Wan and Merwe, 2000).

1. For the forecast step, a particle-based approach is taken.
 - (i) A total of $Y = 2N + 1$ particles, with N equal to the state dimension, are (re)sampled to capture the mean and covariance of the conditional state probability distribution $p[x_{k-1} | \mathcal{Z}_{k-1}]$, by

$$\boldsymbol{\psi}_{k-1|k-1}^i = \begin{cases} \bar{\boldsymbol{\psi}}_{k-1|k-1} & \text{for } i = 1, \\ \bar{\boldsymbol{\psi}}_{k-1|k-1} + \left(\sqrt{(N + \lambda) \cdot P_{k-1|k-1}^x} \right)_i & \text{for } i = 2, \dots, N + 1, \\ \bar{\boldsymbol{\psi}}_{k-1|k-1} - \left(\sqrt{(N + \lambda) \cdot P_{k-1|k-1}^x} \right)_{i-N-1} & \text{for } i = N + 2, \dots, Y, \end{cases} \quad (5.13)$$

where $\lambda = \alpha^2(N + \kappa) - N$ is a scaling parameter, α determines the spread of the particles around the mean and κ is a secondary scaling parameter typically set to 0 (Wan and Merwe, 2000). The vector $\bar{\boldsymbol{\psi}}_{k-1|k-1}$ is the estimated state vector calculated as $\bar{\boldsymbol{\psi}}_{k-1|k-1} = \sum_{i=1}^Y \left(w_{\text{mean}}^i \cdot \boldsymbol{\psi}_{k-1|k-1}^i \right)$, where the weight of each particle’s mean w_{mean}^i and covariance $w_{\text{covariance}}^i$ is given by

$$w_{\text{mean}}^i = \begin{cases} \lambda(N + \lambda)^{-1} & \text{for } i = 1, \\ \frac{1}{2}(N + \lambda)^{-1} & \text{otherwise,} \end{cases}$$

$$w_{\text{covariance}}^i = \begin{cases} \lambda(N + \lambda)^{-1} + (1 - \alpha^2 + \beta) & \text{for } i = 1, \\ \frac{1}{2}(N + \lambda)^{-1} & \text{otherwise,} \end{cases}$$

and β is used to incorporate prior knowledge on the probability distribution. In this chapter, $\beta = 2$ is assumed, which is stated to be optimal for Gaussian distributions (Wan and Merwe, 2000).

- (ii) Each particle is propagated forward in time using the expectation of the nonlinear model, as

$$\begin{aligned} \boldsymbol{\psi}_{k|k-1}^i &= f(\boldsymbol{\psi}_{k-1|k-1}^i, q_{k-1}) & \text{for } i = 1, \dots, Y, \\ \boldsymbol{\zeta}_{k|k-1}^i &= h(\boldsymbol{\psi}_{k|k-1}^i, q_k) & \text{for } i = 1, \dots, Y, \end{aligned} \quad (5.14)$$

where $\boldsymbol{\zeta}_{k|\ell}^i$ is defined as the system output corresponding to the particle $\boldsymbol{\psi}_{k|\ell}^i$.

- (iii) The expected state $\bar{\boldsymbol{\psi}}$ and expected output $\bar{\boldsymbol{\zeta}}$ are calculated as

$$\begin{aligned} \hat{x}_{k|k-1} &= \bar{\boldsymbol{\psi}}_{k|k-1} = \sum_{i=1}^Y \left(w_{\text{mean}}^i \cdot \boldsymbol{\psi}_{k|k-1}^i \right), \\ \hat{z}_{k|k-1} &= \bar{\boldsymbol{\zeta}}_{k|k-1} = \sum_{i=1}^Y \left(w_{\text{mean}}^i \cdot \boldsymbol{\zeta}_{k|k-1}^i \right), \end{aligned} \quad (5.15)$$

and the covariance matrices are estimated from the forecasted ensemble by

$$P_{k|k-1}^x = \sum_{i=1}^Y \left(w_{\text{covariance}}^i \left(\boldsymbol{\psi}_{k|k-1}^i - \bar{\boldsymbol{\psi}}_{k|k-1} \right) \left(\boldsymbol{\psi}_{k|k-1}^i - \bar{\boldsymbol{\psi}}_{k|k-1} \right)^T \right) + Q_{k-1}, \quad (5.16)$$

$$P_{k|k-1}^z = \sum_{i=1}^Y \left(w_{\text{covariance}}^i \left(\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1} \right) \left(\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1} \right)^T \right) + R_k, \quad (5.17)$$

$$P_{k|k-1}^{xz} = \sum_{i=1}^Y \left(w_{\text{covariance}}^i \left(\boldsymbol{\psi}_{k|k-1}^i - \bar{\boldsymbol{\psi}}_{k|k-1} \right) \left(\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1} \right)^T \right) + S_k. \quad (5.18)$$

2. For the analysis step, one can apply the same equations as in Equations 5.10 to 5.12. The UKF has been shown to consistently outperform the ExKF in terms of accuracy, since it uses the nonlinear model for forecasting and covariance propagation. However, this does come at an increased computational cost. Namely, $Y = 2N + 1$ particles are required to capture the mean and covariance of the conditional state probability distribution. This implies that $2N + 1$ function evaluations are required for each UKF update.

Even for a 2-turbine wind farm in WFSim, a computational cost of $1 \cdot 10^2$ s per iteration ($k \rightarrow k+1$) would not be surprising. While Equation 5.14 can easily be parallelized, computational complexity remains troublesome, especially for larger wind farms. The issue of computational complexity is tackled by the Ensemble KF.

5.3.5. The Ensemble Kalman filter (EnKF)

The Ensemble Kalman filter (EnKF) (Evensen, 2003) is very similar to the UKF in that it relies on a finite number of realizations (the “sigma points” or “particles” in the UKF) to approximate the mean and covariance of the conditional probability distribution of x_k knowing \mathcal{Z}_k . However, whereas the UKF relies on a systematic way of distributing the particles such that the mean and covariance of the conditional probability distribution $p[x_k|\mathcal{Z}_k]$ are equal to that of the particles, the EnKF relies on random realizations, without guarantees that the mean and covariance are captured accurately. Though, the EnKF has been shown to work well in a number of applications, with typically far fewer particles than states, i.e., $Y \ll N$ (e.g., Gillijns et al., 2006; Houtekamer and Mitchell, 2005). The forecast and update step are very similar to that of the UKF, namely:

1. In the UKF the particles are redistributed at every timestep, in contrast to the EnKF. Rather, the EnKF propagates the particles forward without redistribution. We define the i^{th} particle as $\boldsymbol{\psi}_{k|\ell}^i \in \mathbb{R}^N$, which is a realization of the conditional probability distribution $p[x_k|\mathcal{Z}_\ell]$. The forecast step is:

- (i) Each particle is propagated forward in time using the nonlinear system dynamics and with the realizations of noise terms w and v denoted by $\hat{w}_{k-1}^i \in \mathbb{R}^N$ and $\hat{v}_k^i \in \mathbb{R}^M$, generated using MATLABs *randn(...)* function.

$$\begin{aligned} \boldsymbol{\psi}_{k|k-1}^i &= f(\boldsymbol{\psi}_{k-1|k-1}^i, q_{k-1}) + \hat{w}_{k-1}^i & \text{for } i = 1, \dots, Y, \\ \boldsymbol{\zeta}_{k|k-1}^i &= h(\boldsymbol{\psi}_{k|k-1}^i, q_k) + \hat{v}_k^i & \text{for } i = 1, \dots, Y. \end{aligned} \quad (5.19)$$

- (ii) The expected state and output are calculated identically as in the UKF using Equation 5.15 with $w_{\text{mean}}^i = (Y-1)^{-1}$. The covariance matrices are estimated from the forecasted ensemble, by

$$P_{k|k-1}^z = \frac{1}{Y-1} \sum_{i=1}^Y \left((\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1}) (\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1})^T \right), \quad (5.20)$$

$$P_{k|k-1}^{xz} = \frac{1}{Y-1} \sum_{i=1}^Y \left((\boldsymbol{\zeta}_{k|k-1}^i - \bar{\boldsymbol{\zeta}}_{k|k-1}) (\boldsymbol{\psi}_{k|k-1}^i - \bar{\boldsymbol{\psi}}_{k|k-1})^T \right). \quad (5.21)$$

2. For the analysis step, one applies Equation 5.10 to determine the Kalman gain L_k . Then, each particle is updated individually, as

$$\boldsymbol{\psi}_{k|k}^i = \boldsymbol{\psi}_{k|k-1}^i + L_k (z_k - \boldsymbol{\zeta}_{k|k-1}^i) \quad \text{for } i = 1, \dots, Y. \quad (5.22)$$

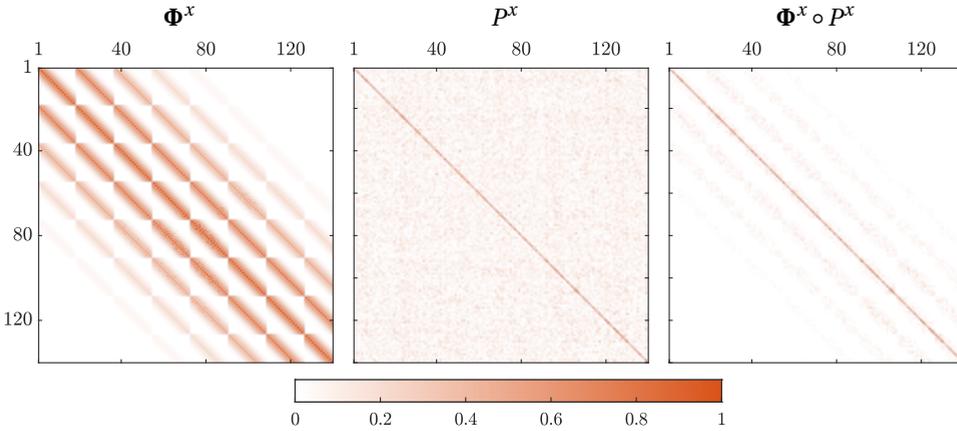


Figure 5.2: Long-range spurious correlations arise in the case where a covariance matrix is described by a small number of particles. Using physical knowledge of the system, these undesired correlations can be corrected. Φ^x is the localization matrix. Applying localization, the covariance of physically nearby states are multiplied with a value close to 1 and the covariance of physically distant states are multiplied with a value close to 0. In our example case, this results in the localized covariance matrix $\Phi^x \circ P^x$, where \circ is the element-wise product.

Note that, in contrast to the ExKF and the UKF, the state covariance matrix P^x (see Equation 5.7 and Equation 5.12) need not be calculated explicitly in the EnKF. This, in combination with the small number of particles $Y \ll N$, is what makes the EnKF computationally superior to the UKF (and often also computationally superior to the ExKF). However, this reduction in computational complexity comes at a price. The disadvantages of the EnKF are discussed in the next section.

5.3.5.1. Challenges in the EnKF for small number of particles

The caveat to representing the conditional state probability distribution with fewer particles than states, $Y \ll N$, is the formation of inbreeding and long-range spurious correlations (Petrie, 2008). The former, inbreeding, is defined as a situation where the state error covariance matrix P^x is consistently underestimated, leading to state estimates that incorrectly rely more on the internal model. One straight-forward method to address this is called “covariance inflation”, in which P^x (or rather, the ensemble from which P^x is calculated) is “inflated” to correct for the underestimated state uncertainty (Petrie, 2008). Mathematically, this is achieved by applying

$$\psi_{k|k-1}^i = \bar{\psi}_{k|k-1} + r \left(\psi_{k|k-1}^i - \bar{\psi}_{k|k-1} \right) \quad \text{for } i = 1, \dots, Y, \quad (5.23)$$

before the analysis step, with $r \in \mathbb{R}$ the inflation factor, typically with a value of 1.01–1.25.

The latter problem, long-range spurious correlations, is visualized in Figure 5.2. In particle-based approaches, the covariance terms cannot be captured exactly. This may lead to the formation of small yet nonzero covariance terms between states and outputs which, in reality, are uncorrelated. This can lead to the drift of unobservable states

and eventually to instability of the KF. Increasing the number of particles is the most straight-forward solution to this problem, but comes at a huge computational cost. A better alternative is “covariance localization”, where physical knowledge of the states and measurements is used to steer the sample-based covariance matrices. Recall that in the surrogate model of Section 5.2, the model states are the velocity and pressure terms inside the wind farm at a physical location. Define that the i^{th} state entry $(x_k)_i$ belongs to a physical location in the farm s_i . Then, looking at an arbitrary state covariance term (i, j) ,

$$\left(P_{k|k-1}^x\right)_{i,j} = \mathbb{E} \left[\left((x_k)_i - (\hat{x}_{k|k-1})_i \right) \left((x_k)_j - (\hat{x}_{k|k-1})_j \right)^T \right],$$

we define the physical distance between these two states as $\Delta s_{i,j} = \|s_i - s_j\|_2$. Now, we introduce a weighting factor into our covariance matrices by multiplying physically distant states with a value close to 0 and multiplying physically nearby states with a value close to 1. A popular choice for such a weighting function is Gaspari-Cohn’s fifth-order discretization of a Gaussian distribution (Gaspari and Cohn, 1999), given by

$$\phi(c_{i,j}) = \begin{cases} -\frac{1}{4}c_{i,j}^5 + \frac{1}{2}c_{i,j}^4 + \frac{5}{8}c_{i,j}^3 - \frac{5}{3}c_{i,j}^2 + 1 & \text{if } 0 \leq c_{i,j} \leq 1, \\ \frac{1}{12}c_{i,j}^5 - \frac{1}{2}c_{i,j}^4 + \frac{5}{8}c_{i,j}^3 + \frac{5}{3}c_{i,j}^2 - 5c_{i,j} + 4 - \frac{2}{3}\frac{1}{c_{i,j}} & \text{if } 1 < c_{i,j} \leq 2, \\ 0 & \text{otherwise,} \end{cases} \quad (5.24)$$

with $c_{i,j} = \frac{\|\Delta s_{i,j}\|_2}{L}$ a normalized distance measure, with L the cut-off distance. Applying Equation 5.24 for the covariance matrices $P_{k|k-1}^z$ and $P_{k|k-1}^{xz}$, we can define the localization matrices

$$\Phi^z = \begin{bmatrix} \phi(c_{1,1}^z) & \cdots & \cdots \phi(c_{1,M}^z) \\ \vdots & \ddots & \\ \phi(c_{M,1}^z) & & \phi(c_{M,M}^z) \end{bmatrix}, \quad \Phi^{xz} = \begin{bmatrix} \phi(c_{1,1}^{xz}) & \cdots & \cdots \phi(c_{1,M}^{xz}) \\ \vdots & \ddots & \\ \phi(c_{N,1}^{xz}) & & \phi(c_{N,M}^{xz}) \end{bmatrix},$$

where $c_{i,j}^z$ is the normalized distance between two measurements i and j and $c_{i,j}^{xz}$ is the normalized distance between state i and measurement j , respectively. Finally, localization and inflation can be incorporated into Equations 5.20 and 5.21 by

$$P_{k|k-1}^z = \Phi^z \circ \frac{1}{Y-1} \sum_{i=1}^Y \left(\left(\zeta_{k|k-1}^i - \bar{\zeta}_{k|k-1} \right) \left(\zeta_{k|k-1}^i - \bar{\zeta}_{k|k-1} \right)^T \right), \quad (5.25)$$

$$P_{k|k-1}^{xz} = r \cdot \Phi^{xz} \circ \frac{1}{Y-1} \sum_{i=1}^Y \left(\left(\zeta_{k|k-1}^i - \bar{\zeta}_{k|k-1} \right) \left(\psi_{k|k-1}^i - \bar{\psi}_{k|k-1} \right)^T \right), \quad (5.26)$$

where \circ is the element-wise product (Hadamard) of the two matrices. The improvement in terms of computational efficiency and estimation performance is displayed in Figure 5.3. A significant increase in performance is shown, especially for smaller numbers of particles. This is in agreement with what was seen in previous work (Doekemeijer et al., 2017). Furthermore, performance is more consistent. Additionally, note that there is no increase in computational cost, as the covariance matrices are made sparse, leading to a cost reduction in the calculation of Equation 5.10, which makes up for the extra

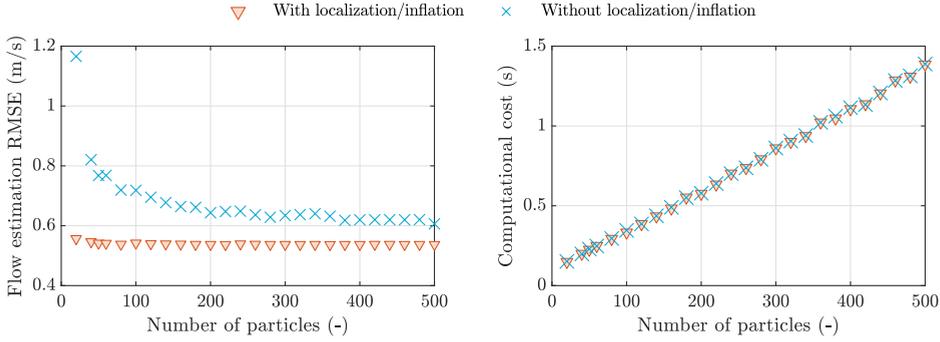


Figure 5.3: This figure shows the estimation performance and computational cost (parallelized, 8 cores) of the EnKF for a range of ensemble sizes, with and without inflation and localization. Great improvement is seen for estimation accuracy, at no additional computational cost. The simulation scenario is described in detail in Section 5.4.2 and the results presented here are rather meant as an indication.

operations of Equations 5.25 and 5.26. Also, note that the localization matrices are time-invariant and can be calculated offline.

5.3.6. Synthesizing an online model calibration solution

Certain model parameters such as ℓ_s are closely related to the turbulence intensity, which vary over time. Estimation of such parameters is achieved by extending the state vector with (a subset of) the model parameters. In this chapter, ℓ_s is concatenated to the state vector as random walk model, with a certain standard deviation (covariance). Higher values of ℓ_s lead to more wake recovery, making the calibration solution adaptable to varying turbulence levels. This adds one scalar entry to x_k , which is a negligible addition in terms of computational cost.

Furthermore, a proposal is made for the estimation of the freestream wind speed U_∞ . This is suggested to be done using the turbine power generation measurements, following the ideas of Gebraad et al. (2016) and Shapiro et al. (2017b). Using the wind vanes and employing a simple steady-state wake model from the literature (Mittelmeier et al., 2017), the turbines operating in freestream flow can be distinguished from the ones operating in waked flow. Next, define $\Gamma \in \mathbb{Z}^{\aleph}$ as a vector specifying the upstream turbines, with \aleph the total number of turbines operating in freestream. Then, the instantaneous rotor-averaged flow speed at each turbine hub can be estimated by inverting the turbine power expression from WFSim (Boersma et al., 2018a). One wind-farm-wide freestream wind speed U_∞ is then calculated using actuator disk theory. Smoothing results with a low-pass filter with time-constant c_{u_∞} on the average of $U_{\infty i}$ for each upstream turbine i , we obtain

$$c_{u_\infty} \frac{\partial U_\infty}{\partial t} = \frac{1}{\aleph} \sum_{i \in \Gamma} \left(\sqrt[3]{\frac{P_{\text{turb},i}^{\text{measurements}}}{\frac{c_p}{2} \rho A C'_{T_i} \cos(\gamma_i)^3} \cdot \left(1 + \frac{1}{4} C'_{T_i}\right)} \right) - U_\infty, \quad (5.27)$$

where it is assumed that $U_{\infty_i} \approx U_{r_i} \left(1 + \frac{1}{4} \cdot C'_{T_i}\right)$ when $\gamma_i \approx 0$, with U_{r_i} the wind speed at the rotor of turbine i . Moreover, c_p is a tuning factor and C'_T is the turbine control input representative of the generator torque and blade pitch angles, as previously discussed in Section 5.2. Research is currently ongoing on how to best incorporate the effects of turbine yaw ($\gamma \neq 0$) into the definition of C'_T . Furthermore, ρ is the air density, A is the rotor swept area and $P_{\text{turb},i}^{\text{measurements}}$ is the measured instantaneous power capture of turbine i .²

Combining these elements yields an efficient, modular and accurate model calibration solution for WFSim. The model states are estimated using SCADA and/or LiDAR data, of which the former is readily available and the latter becoming more popular. State estimation paired with parameter estimation improves the accuracy of the surrogate model, potentially leading to more accurate control. Additionally, the freestream wind speed is estimated using readily available SCADA data. This control solution is implemented in MATLAB and leverages the numerically efficient precompiled solvers and parallelization for model propagation. The EnKF is orders of magnitude faster than existing estimation algorithms due to covariance localization and inflation, while competing with the UKF in terms of accuracy.

5.4. Results

In this section, the calibration solution detailed in Section 5.3 will be validated using high-fidelity simulations. First, the model used to generate the validation data will be described in Section 5.4.1. Then, simulation results for a two-turbine and a nine-turbine case are presented in Sections 5.4.2 and 5.4.3, respectively.

Note that for the presented results, pressure terms are ignored in the state vector, as they appeared unnecessary for state estimation in previous work (Doekemeijer et al., 2017). Furthermore, for simplicity and due to lack of information, the process and measurement noise will be assumed to be uncorrelated, $S_k = 0$ and Q_k and R_k are assumed to be time-invariant and diagonal. Also, note that the simulations presented are not conclusive on the feasibility of the solution under all relevant conditions experienced in an operational wind farm. Rather, this chapter presents a first step towards algorithm validation.

5.4.1. SOWFA

High-fidelity simulation data is generated using the Simulator fOr Wind Farm Applications (SOWFA), developed by the National Renewable Energy Laboratory. SOWFA provides accurate flow data at a fraction of the cost of field tests. It solves the filtered, three-dimensional, unsteady, incompressible Navier-Stokes equations over a finite temporal and spatial mesh, accounting for the Coriolis and geostrophic forcing terms. SOWFA is a large-eddy simulation solver, meaning that larger scale dynamics are resolved directly and turbulent structures smaller than the discretization are approximated using

²Note that this method for the estimation of U_{∞} relies solely on power measurements and therefore only works for below-rated conditions. For estimation of U_{∞} in above-rated conditions, one may require the implementation of a wind speed estimator on each turbine (e.g., Simley and Pao (2016)).

Table 5.1: Overview of several settings for the SOWFA and the WFSim 2-turbine wind farm simulation.

Variable	SOWFA	WFSim
Domain size	3.0km \times 3.0km \times 1.0km	1.9km \times 0.80km
Cell size near rotors	3m \times 3m \times 3m	38m \times 33m
Cell size outer regions	12m \times 12m \times 12m	38m \times 33m
Rotor model	ALM	ADM ($c_f = 1.4$, $c_p = 0.95$)
Inflow wind speed	8.0 m/s	8.0 m/s
Atmospheric turbulence	Low, $TI_\infty = 5.0\%$	$d' = 1.8 \cdot 10^2$ m, $d = 6.1 \cdot 10^2$ m, $\ell_s = 1.8 \cdot 10^{-2}$

subgrid-scale models to suppress computational cost (Churchfield et al., 2012a). The turbine rotor is modeled using an actuator line model (ALM) as derived from Sørensen and Shen (2002). SOWFA has previously been used for lower-fidelity model validation, controller testing and to study the aerodynamics in wind farms (e.g., Fleming et al., 2016, 2017a; Gebraad et al., 2017). The interested reader is referred to Churchfield et al. (2012a) for a more in-depth description of SOWFA and LES solvers in general.

5.4.2. 2-turbine ALM with turbulent inflow

In this section, a two-turbine wind farm is simulated to analyze the effect of different measurement sources, KF algorithms and the difference between state-only and state-parameter estimation. This simple wind farm contains two NREL 5-MW baseline turbines with $D = 126.4$ m, separated $5D$ in stream-wise direction. This LES simulation was described in more detail in Annoni et al. (2016b). Important simulation properties are listed in Table 5.1 for SOWFA and WFSim. The effect of the turbulence intensity on the wake dynamics in SOWFA is captured in WFSim through its mixing-length turbulence model. In these simulations, WFSim is purposely run with a too low value for ℓ_s in order to represent the realistic situation of a model mismatch. The remaining tuning parameters in WFSim were chosen such that a weighted-sum cost function of the power and flow errors was minimized.

Firstly, the three KF variants will be compared in Section 5.4.2.1. Secondly, in Section 5.4.2.2, estimation using different information sources is compared. Thirdly, the potential of joint state-parameter estimation is displayed in Section 5.4.2.3.

5.4.2.1. A comparison of the KF variants for state estimation

In this simulation study, four estimation cases are compared: 1) the ExKF, 2) the UKF, 3) the EnKF and 4) the open-loop (OL) simulation, i.e., without estimation. The focus here is on state-only estimation, thus excluding ℓ_s . Flow measurements downstream of each turbine are assumed (e.g., using LiDAR), their locations denoted as red dots in Figure 5.4, which is about 2% of the full to-be-estimated state space. These measurements are artificially disturbed by zero-mean white noise with $\sigma = 0.10$ m/s. The KF settings are listed in Tables 5.2 and 5.3. The KF covariance matrices were obtained through an iterative tuning process in previous work (Doekemeijer et al., 2017) with minor adjustments,

Table 5.2: Covariance settings for the KF variants, with I , the $\mathbb{R}^{* \times *}$ identity matrix. The full covariance matrices are diagonal concatenations of the entries. For example, P_0 is $\text{diag}(P_{0,u}, P_{0,v})$ and $\text{diag}(P_{0,u}, P_{0,v}, P_{0,\ell_s})$ for state-only and state-parameter estimation, respectively.

Variable	Symbol	Units	Value
Initial state error covariance of \vec{u}_k	$P_{0,u}$	$(\text{m/s})^2$	$1.0 \cdot 10^{-1} \cdot I_{N_u}$
Initial state error covariance of \vec{v}_k	$P_{0,v}$	$(\text{m/s})^2$	$1.0 \cdot 10^{-1} \cdot I_{N_v}$
Initial state error covariance of ℓ_{s_k}	P_{0,ℓ_s}	–	$5.0 \cdot 10^{-1}$
Model error covariance of \vec{u}_k	$Q_{0,u}$	$(\text{m/s})^2$	$1.0 \cdot 10^{-2} \cdot I_{N_u}$
Model error covariance of \vec{v}_k	$Q_{0,v}$	$(\text{m/s})^2$	$1.0 \cdot 10^{-4} \cdot I_{N_v}$
Model error covariance of ℓ_{s_k}	Q_{0,ℓ_s}	–	$1.0 \cdot 10^{-4}$
Measurements error covariance of flow	$R_{u,v}$	$(\text{m/s})^2$	$1.0 \cdot 10^{-2} \cdot I_{M_{u,v}}$
Measurements error covariance of P	R_P	$(\text{W})^2$	$1.0 \cdot 10^8 \cdot I_{N_T}$

Table 5.3: Choice of tuning parameters for the KF variants, for both the 2-turbine and 9-turbine simulation cases. Note that the ExKF does not support power measurements nor parameter estimation due to the lack of linearization and does not have any additional tuning parameters. In terms of computational cost: simulations were run on a single node using 8 cores in parallel.

	2-turbine		2-turbine		2-turbine		9-turbine	
Variable	ExKF	UKF	EnKF	EnKF	EnKF	EnKF	EnKF	EnKF
Number of particles, Y	–	4275	50	50	50	50	50	50
Tuning parameters	–	α 1.0 β 2.0 κ 0	L 131 m r 1.025					
Computational cost/iteration	16.2 s	14.0 s	0.25 s	0.25 s	0.25 s	0.25 s	1.2 s	1.2 s

to simulate performance for untrained data. Figure 5.4 shows state (flow field) estimation of the three KF variants for two time instants, $t = 300$ s and $t = 700$ s. In this figure, $(\Delta \vec{u})_{\bullet} \in \mathbb{R}^{N_u}$ is defined as the absolute error between the estimated and true longitudinal flow velocities in the field.

Looking at Figure 5.4, the open-loop estimations are accurate for the unwaked and single waked flow, yet are lacking in the situation of two overlapping wakes, for which the KFs correct. There is no significant difference in accuracy between the different KF variants, yet they differ by two orders of magnitude in computational cost (Table 5.3).

5.4.2.2. A comparison of sensor configurations

Previous results (Doekemeijer et al., 2016, 2017) have relied on flow measurements for state estimation. However, in existing wind farms, such measurements are typically not available. Rather, readily available SCADA data should be used for the purpose of model calibration. For this reason, state estimation with the EnKF leveraging instantaneous turbine power measurements, using an upstream-pointing LiDAR and using a downstream-pointing LiDAR are compared in Figure 5.5. Flow and power measurements are artificially disturbed by zero-mean white Gaussian noise with $\sigma = 0.10$ m/s and $\sigma = 10^4$ W, respectively.

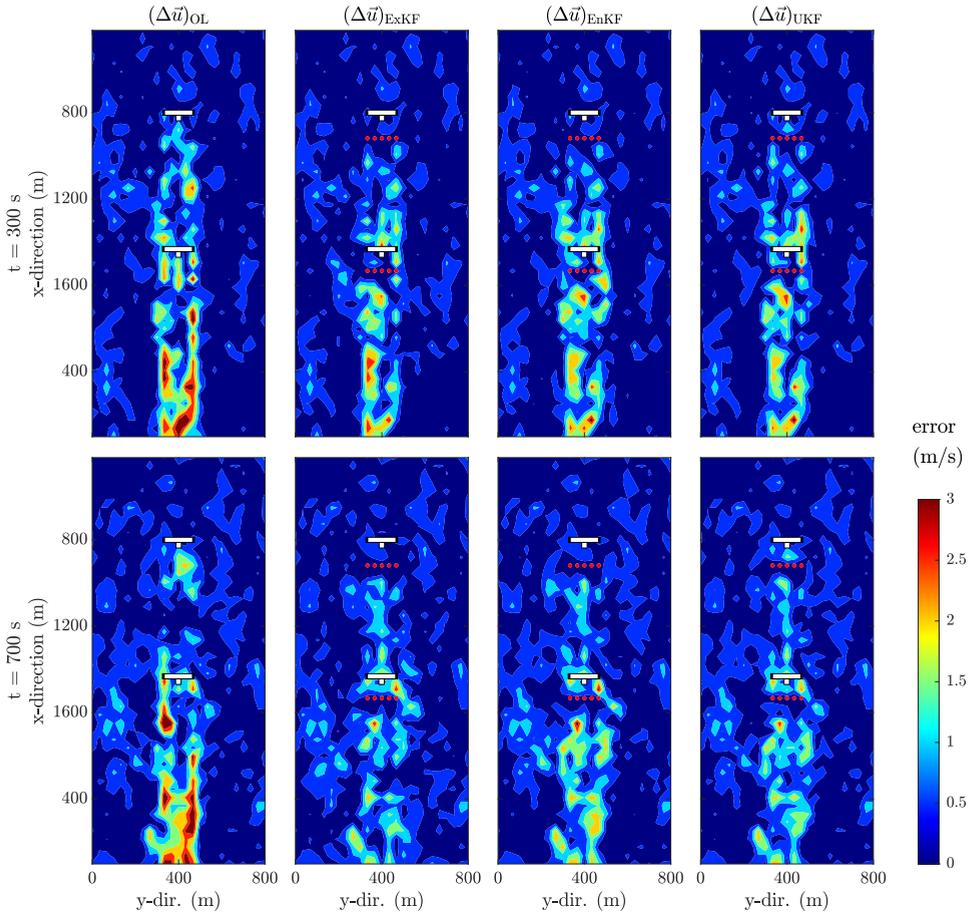


Figure 5.4: Comparison of absolute values of the estimation errors (in longitudinal flow fields) for state-only estimation with the ExKF, EnKF and UKF at $t = 300$ s and $t = 700$ s, with $(\Delta \vec{u})_{\bullet} = |\vec{u}_{\bullet} - \vec{u}_{SOWFA}|$. The model and KF settings are depicted in Tables 5.1, 5.2 and 5.3. Wind is coming in from the top, flowing towards the bottom. The measured states are depicted by red dots in the flow, not to be confused with estimation error. The KFs consistently improve the instantaneous flow field estimations, noticeably nearby the measurements.

The KF settings are displayed in Tables 5.2 and 5.3. In Figure 5.5 it can be seen that SCADA data allows comparable performance compared to the use of flow measurements, making the proposed closed-loop control solution feasible for implementation in existing wind farms, without the need for additional equipment. Furthermore, this modular framework allows the use of a combination of LiDAR systems, measurement towers and/or SCADA data, whichever is available, for model calibration.

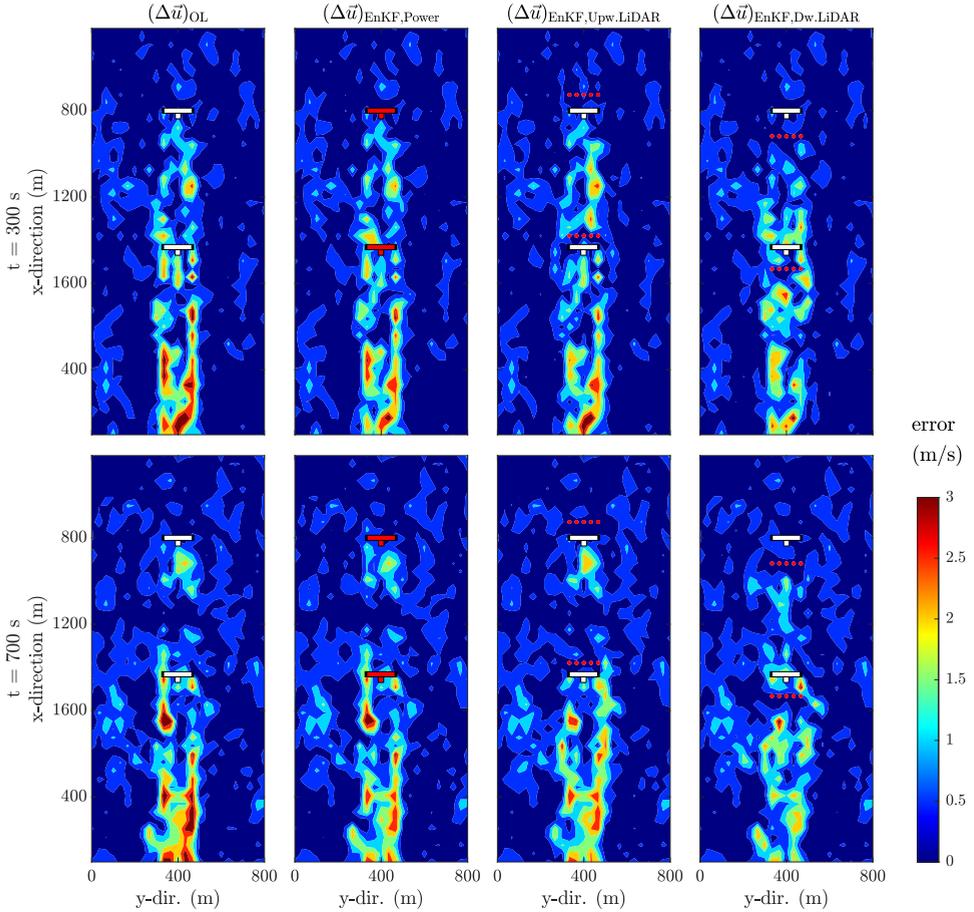


Figure 5.5: Comparison of absolute values of the estimation errors (in longitudinal flow fields) for state-only estimation with the EnKF for various sensor configurations: using turbine power measurements, using flow measurements with a LiDAR system pointing upstream and using flow measurements with a LiDAR system pointing downstream of the rotor. Here, $(\Delta \vec{u})_{\bullet} = |\vec{u}_{\bullet} - \vec{u}_{SOWFA}|$. Here, wind flows from top to bottom. The sensors are depicted by red dots (flow measurements) or red turbines (power measurements), not to be confused with estimation error.

5.4.2.3. Joint state-parameter estimation

Forecasting, as used in predictive control, benefits from the calibration of model parameters in addition to the states. Joint state-parameter estimation using flow measurements downstream of each turbine (as shown in the rightmost plots in Figure 5.5) disturbed by zero-mean white noise with $\sigma = 0.10$ m/s for the EnKF and UKF is displayed in Figure 5.6. The KF settings are shown in Tables 5.2 and 5.3. At $t = 0$ s, both the OL and the KF simulations start with the same (wrong) value for ℓ_s . Then, every second,

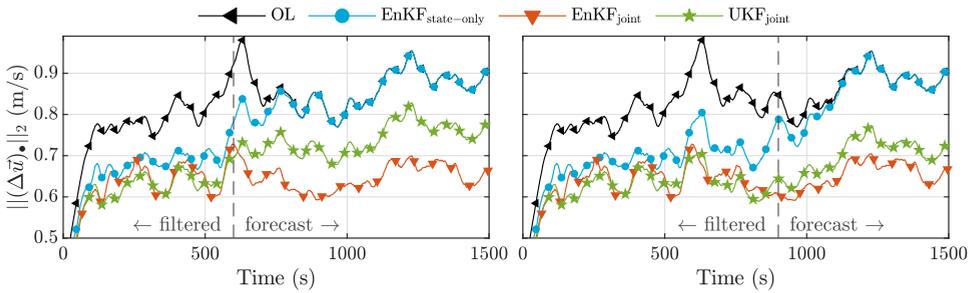


Figure 5.6: Comparison of forecasting performance for state-only and joint state-parameter (ℓ_s) estimation with the EnKF and UKF, where measurements are available up until the vertical dashed lines, after which the estimation becomes a forecast. Here, the 2-norm of the estimation error is plotted along the y -axis, with $(\Delta \bar{u})_\bullet = |\bar{u}_\bullet - \bar{u}_{\text{SOWFA}}|$.

(noisy) measurements are fed into the KFs and the state vector as well as the model parameter ℓ_s are estimated. However, for the OL simulation, no measurements are fed in: the state vector is simply updated with the nominal model and the value for ℓ_s remains the same throughout the simulation. Now, after 600 s (left plot in Figure 5.6) and 900 s (right plot in Figure 5.6), a forecast is started, meaning no measurements are available after that time. At that moment, the OL model still has the same (poor) value for ℓ_s as at $t = 0$ s, while the value for ℓ_s in the KFs has improved. From Figure 5.6, it becomes clear that the estimates are not only improved for the 3-minute forecast, but are also consistently better than the non-calibrated (open-loop) model's 10-minute forecast due to the estimation of ℓ_s .³ Furthermore, the EnKF performs comparably to the UKF at a lower computational cost. Note that the EnKF even outperforms the UKF in this simulation, expected to be due to randomness in the EnKF.

5.4.3. 9-turbine ALM with turbulent inflow

In this section, we investigate the performance of the EnKF-based model calibration solution under a more realistic 9-turbine wind farm scenario. The purpose of this case study is to highlight the need for state-parameter estimation for accurate wind farm modeling. The wind farm contains nine NREL 5-MW turbines, oriented in a three by three layout, separated $5D$ and $3D$ in stream- and cross-wise direction, respectively. The turbines start with a 30° yaw misalignment, but are then aligned with the mean wind direction within the first 30 s of simulation. The turbine layout and numbering is shown in the top-left subplot of Figure 5.8. This LES simulation has been used before in the literature and is described in more detail in Boersma et al. (2018a). A number of important simulation properties are listed in Table 5.4 for SOWFA and WFSim, respectively.

Compared to the 2-turbine case, N has increased by a factor 4. In the UKF, this would result in the same factor of additional particles. Thus, not only is each particle more expensive to calculate, there are also more particles. Rather, in the EnKF, the approach

³Note that this is highly dependent on the frequency at which the freestream conditions change in the atmosphere.

Table 5.4: Overview of several settings for the SOWFA and the WFSim 9-turbine wind farm simulation.

Variable	SOWFA	WFSim
Domain size	3.5km × 3.0km × 1.0km	1.9km × 0.80km
Cell size near rotors	3m × 3m × 3m	25m × 38m
Cell size outer regions	12m × 12m × 12m	25m × 38m
Rotor model	ALM	ADM ($c_f = 2.0$, $c_p = 0.97$)
Inflow wind speed	12.03 m/s	9.0 m/s (OL) and 12.0 m/s (OL) 9.0 m/s (EnKF)
Atmospheric turbulence	Low, $TI_\infty = 4.7\%$	$d' = 3.8 \cdot 10^1$ m, $d = 5.2 \cdot 10^2$ m, $\ell_s = 3.9 \cdot 10^{-2}$

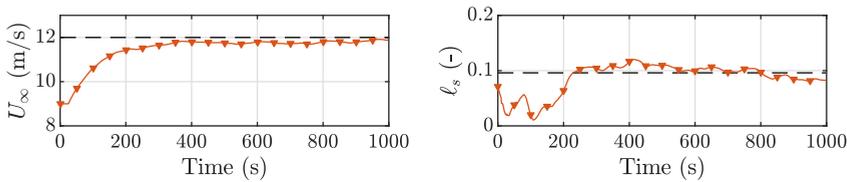


Figure 5.7: Convergence of ℓ_s and U_∞ using the EnKF. In dashed lines are the grid-searched optimal constant values for the open-loop simulation. With power measurements only, the EnKF is able to estimate these parameters successfully in addition to the model states.

is heuristic. None of the EnKF settings needed to be changed for good performance compared to Section 5.4.2, as displayed in Tables 5.2 and 5.3.

As shown in Table 5.3, the EnKF has a low computational cost of 1.2 s/iteration (8 cores, parallel). In this case study, both the complete model state (flow field), the turbulence model parameter ℓ_s and the freestream flow speed U_∞ are estimated in real-time using exclusively (readily available) power measurements from the turbines. The EnKF and one of the open-loop simulations (OL) will deliberately be initialized with a poor value for ℓ_s and U_∞ to investigate convergence. The other open-loop simulation will be initialized with a poor value for ℓ_s but a correct value for U_∞ for comparison.

In Figure 5.7, it can be seen that the EnKF is successful in estimating U_∞ and ℓ_s after 300 s using only wind turbine power measurements. Furthermore, the flow fields of SOWFA, of the open-loop (OL) simulation with $U_\infty = 9.0$ m/s and of the EnKF at various time instants are displayed in Figure 5.8. From this figure, it can be seen that the EnKF has large errors at the start of the simulation. However, after 10 s, the error in flow states surrounding each turbine significantly decreases through the use of turbine power measurements. This estimated flow then propagates downstream, “clearing up” the errors in the vicinity of the wind turbines. As time further propagates, the freestream estimation improves and finally the estimation error converges.

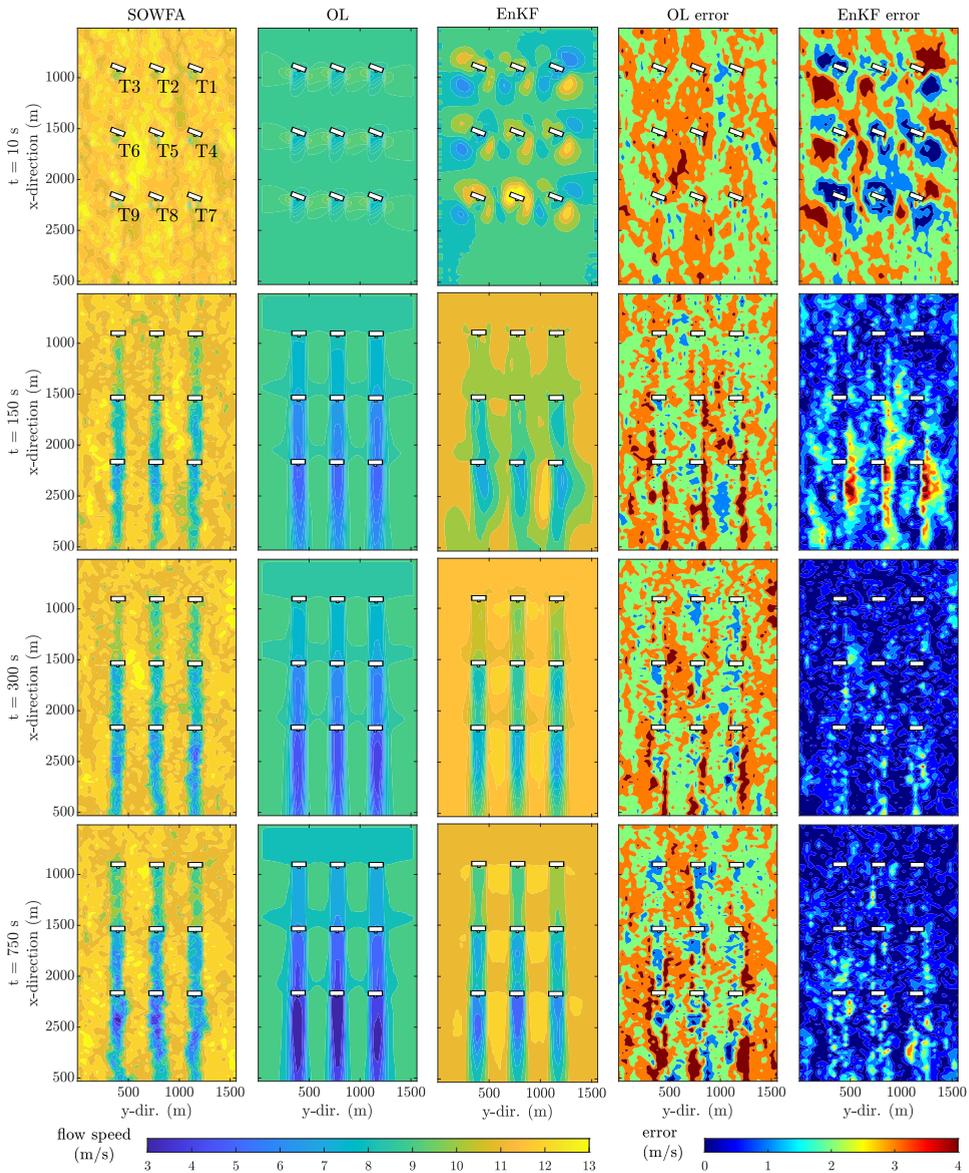


Figure 5.8: Comparison of absolute values of the estimation errors (in longitudinal flow fields) for state-parameter estimation with the EnKF. Wind is coming in from the top and flows downwards. The variables U_∞ and ℓ_s are incorrectly initialized in both the OL and the EnKF. In the EnKF, U_∞ and ℓ_s are estimated in addition to the states, using only turbine power measurements. The EnKF quickly converges for the states near the turbines as seen at $t = 10$ s and more slowly at the outer states. This aligns with the slow convergence of ℓ_s and U_∞ seen in Figure 5.7. After 300 s, the EnKF has converged to a negligible estimation error.

Table 5.5: Turbine-averaged RMSE in power timeseries of Figure 5.9 (compared to SOWFA). The lower the RMSE, the better the forecast.

turbine row	OL ($U_\infty = 9.0$ m/s)	OL ($U_\infty = 12.0$ m/s)	EnKF
1	1.46 MW	0.19 MW	0.16 MW
2	1.61 MW	0.30 MW	0.18 MW
3	1.78 MW	0.82 MW	0.32 MW

The power forecasting performance is shown in Figure 5.9 and Table 5.5. As also seen in Figure 5.7, the EnKF converges after 300 s and indeed the power forecasts outperform those of the OL simulation at $t = 300$ s. Furthermore, it is interesting to see that the filtered power estimates of the first row of turbines ($i = 1, 2, 3$) starts low at $t = 1$ s, but converges to the true power at $t \approx 200$ s. This can be related to the mismatch in U_∞ , which takes approximately 300 s to converge to the true value of 12 m/s, as seen in Figure 5.7. The oscillatory behavior in both the OL and EnKF power predictions is due to the absence of rotor inertia in the rotor model, turbulent structures in the flow and large fluctuations on the excitation signal C'_T .

Finally, the forecasts for flow at times $t = 300$ s and $t = 600$ s are examined in Figure 5.10. The large flow estimation mismatch in the EnKF at $t < 250$ s quickly reduces and for $t \geq 250$ s the EnKF estimation is consistently better than both the OL cases. This has to do with the convergence of the model parameters ℓ_s and U_∞ and the estimation of the states surrounding the turbines using the power measurements.

A crucial remark with the simulations presented here is that low-frequency changes in the atmosphere are neglected. In a real wind farm, atmospheric properties such as the mean wind direction and turbulence intensity change continuously and this will impact the estimation and forecasting performance. The EnKF uses an assumption of persistence for the atmospheric properties at the time of forecasting and thus a change in mean wind direction may invalidate the model forecast.

5.5. Conclusions

This chapter presented a real-time model calibration algorithm for the dynamic wind farm model “WFSim”, relying on an Ensemble Kalman filter (EnKF) at its core. The joint state-parameter calibration solution was tested in two high-fidelity simulation case studies. Using exclusively SCADA measurements which are readily available in current wind farms, the adaptability to model discrepancies in a 9-turbine wind farm simulation was shown, at a low computational cost of 1.2 s per timestep on an 8-core CPU. Specifically, the freestream wind speed and turbulence intensity were shown to converge to their optimal values within 300 s. Furthermore, the EnKF was shown to perform comparably in terms of accuracy to the state-of-the-art algorithms in the literature, at a computational cost of multiple orders of magnitude lower. Additionally, estimation using flow measurements from LiDAR was compared to estimation using SCADA data and it was shown that SCADA data can effectively be used for real-time model calibration. In future work, the algorithm presented here should be tested under high-fidelity simula-

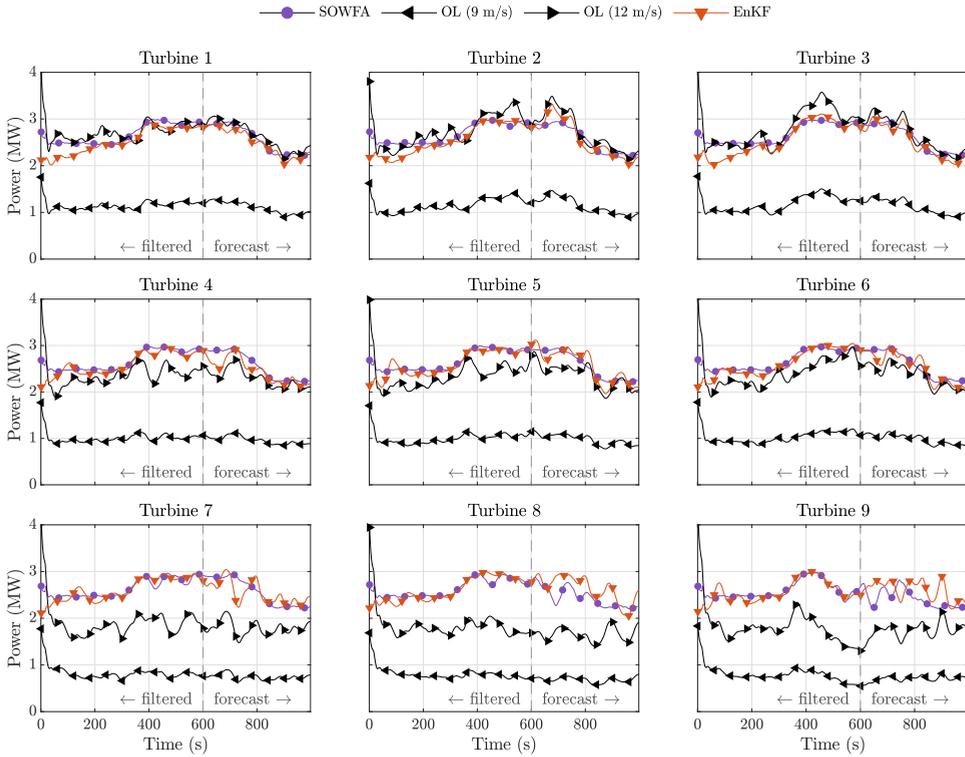


Figure 5.9: Comparison of power forecasting using the EnKF with measurements available up until time $t = 600$ s. The convergence of U_∞ is seen as a positive slope in the power signal until $t = 250$ s in the first row of turbines (aligning with Figure 5.7). After convergence, forecasting is better than in open-loop. Oscillatory behavior is still present due to an oscillatory input signal (C'_T), turbulent flow field and the absence of inertia in the rotor model.

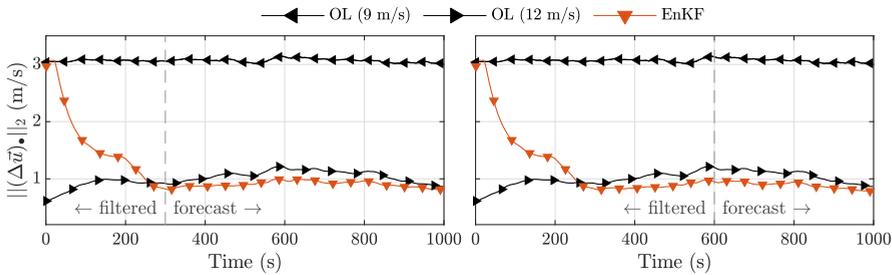


Figure 5.10: Comparison of flow field estimation for the 9-turbine case. Measurements are available until $t = 300$ s (left) and $t = 600$ s (right), respectively. The EnKF converges to the true U_∞ after 300 s. After convergence, the forecasts are significantly better than in open-loop.

tions with realistic low-frequency changes. This would provide insight into the potential of the work at hand and advance towards a practical wind farm implementation. This chapter presented an essential building block for real-time closed-loop wind farm control using surrogate dynamic wind farm models.

6

CONCLUSIONS AND RECOMMENDATIONS

ABSTRACT

In the introduction of this thesis, the overarching objective of this dissertation was formalized as “Maturing wind farm control technologies for power maximization in a model-based closed-loop framework towards real-world practical applicability.” In conclusion, this dissertation has matured the state of the art on various fronts. First, steady-state models were shown to be insufficient in the current wind farm control solutions. Accordingly, a model adaptation solution was developed to increase model accuracy using measurements from the wind farm in real time, shifting towards the closed-loop paradigm. High-fidelity simulations showed the resilience of the proposed closed-loop wind farm controller to time-varying inflow conditions, but also suggested there is room for improvement by addressing the temporal dynamics in the control solution. Correspondingly, this thesis proposed dynamic surrogate wind farm models as a substitute for steady-state models. A real-time model adaptation solution was proposed for a medium-fidelity dynamic surrogate wind farm model at a low computational cost. This estimation algorithm was validated in high-fidelity simulation, successfully reconstructing the second-to-second two-dimensional flow field. At large, closed-loop model-based wind farm control is paving the way for practicable wake steering. The contributions in this dissertation greatly advance the status quo of wind farm control solutions and their applicability in commercial wind farms.

6.1. Conclusions

The overarching objective of this thesis was formulated in Section 1.8 as “Maturing wind farm control technologies for power maximization in a model-based closed-loop framework towards real-world practical applicability”.

At large, this dissertation has successfully matured the steady-state-model-based wind farm control concept on various fronts, essentially addressing the issue of model uncertainty. Closed-loop model-based wind farm control is paving the way for practicable wake steering. The contributions in this dissertation greatly advance the status quo of wind farm control solutions and their applicability in commercial wind farms.

More precisely, three research questions were derived from the overarching objective in Section 1.8. Consequently, several scientific contributions were presented throughout this dissertation to answer these questions.

Research question I: “What are the limitations of the current open-loop wind farm control solutions?”

Chapter 2 presented the results from a field experiment demonstrating wake steering at a commercial onshore wind farm at the coast of Italy. The algorithm leveraged therein uses the steady-state surrogate wind farm model FLORIS in an open-loop wind farm control framework. In the field experiment, two- and three-turbine interaction were looked into, where the upstream one or two turbines were misaligned with the inflow.

The results showed that wake steering has potential to increase the power production significantly, reporting gains of up to 35% for two-turbine interaction and up to 16% for three-turbine interactions. However, unnecessary or erroneous wake steering were seen for various wind directions too, leading to performance losses. In addition to the gains achieved through wake steering at downstream turbines, more interesting to note is that a significant share in gains were from the upstream turbines, showing an increased power production of the yawed turbine itself compared to baseline operation for some wind directions. The surrogate wind farm model, while capturing the general trends of wake interaction, lacked the details necessary to accurately represent the measurements. While the surrogate model assumes a flat terrain, in reality the wind farm is located in complex terrain, with an average elevation of 360 m to 400 m above sea level, surrounded by hills of 400 m to 450 m above sea level. This complex terrain and the influence it has on the airflow are ignored in the FLORIS model. Moreover, some of the wind turbines in the wind farm have a different hub height. While FLORIS does capture the change in absolute location of the rotor, the changes in the downstream flow and its interaction with downstream turbines are not captured explicitly and it remains uncertain whether FLORIS can accurately capture these effects. Finally, FLORIS ignores the inherent time-varying nature of the flow and the wind turbines since it is a steady-state model. These assumptions are expected to have fundamentally compromised the reliability of the FLORIS model.

This field experiment showed that the proposed open-loop wind farm control solution does not consistently achieve energy gains through wake steering at this complicated, onshore wind farm. This thesis therefore pushes the state of the art on two

fronts: establishing resilience to model uncertainty in steady-state-model-based wind farm control solutions and pushing for the usage of dynamic surrogate models in wind farm control solutions.

Research question II: “Can we establish resilience to surrogate model uncertainty in modern wind farm control solutions through real-time measurement feedback and if so, how?”

In view of Chapter 2, it was apparent that the current wind farm control solutions do not suffice, hypothesized to be due to model inaccuracies. Accordingly, Chapters 3 and 4 focused on transitioning to the closed-loop wind farm control architecture – using measurements to estimate the inflow wind conditions and adapt the surrogate wind farm model in real time to accurately capture the current wind farm behavior.

The theoretical analysis presented in Chapter 3 shed light on the measurements that are necessary to accurately reconstruct the ambient conditions using steady-state wind farm models. The theoretical framework confirmed the notion that local wind speed and wind direction estimates at each turbine are requisites for the estimation of a wind-farm-wind wind speed and wind direction, respectively. Additionally, it became apparent that wake interaction is necessary to derive the turbulence intensity from the measurements in a wind farm, assuming no other sensors are available. This conflicts with the objective of wake steering, in which wake interactions are minimized. Moreover, the analysis showed that more complicated, unstructured wind farm layouts yield a higher observability than structured layouts due to the increase in unique wake interaction. Inferring from this theoretical analysis, it became clear that the degree of observability of the ambient conditions in wind farms is inconsistent. Therefore, estimation algorithms require careful implementation, taking this (lack of) observability into account.

Correspondingly, a closed-loop steady-state-model-based wind farm control solution was devised in Chapter 4. The estimation algorithm within this framework uses the theoretical measure of observability to ensure correct estimation of the ambient conditions. This wind farm control algorithm was then tested in high-fidelity simulation, subjected to time-varying inflow conditions, as the first of its kind in the literature. The wake steering algorithm led to an averaged 1.4% increase in power production for a 6-turbine wind farm. This simulation showed the resilience of the proposed controller to more realistic, time-varying inflow conditions and thereby solidified the control solution as the first realistic, closed-loop wind farm control solution for yaw-based wake steering.

While the net gain resulting from the proposed closed-loop wind farm controller was positive, there were time periods in which the power production was several percents lower than in traditional, locally greedy operation. These losses mainly occurred when the inflow conditions were changing in time, leading to larger model mismatches and erroneous turbine misalignment. To address this issue, temporal dynamics should be incorporated into the control solution. Fundamentally, this implies that the focus should be shifted from steady-state towards dynamic surrogate wind farm models.

Research question III: “What benefits do dynamic surrogate models have for wind farm control and how can they be implemented?”

The performance losses in the high-fidelity simulations of Chapter 4 occurred shortly after the inflow wind properties had changed and the solution therefore lies in the usage of dynamic surrogate models in wind farm control solutions. However, dynamic models often go paired with a significant increase in computational cost and algorithm complexity, which have prevented the adoption of such models in practice.

Chapter 5 focused on the development of a real-time model adaptation (state estimation) algorithm for the nonlinear dynamic surrogate wind farm model WFSim, relying on an Ensemble Kalman filter at its core. High-fidelity simulation showed that the algorithm is able to accurately reconstruct and track the second-to-second two-dimensional flow field at turbine hub height in a virtual 9-turbine offshore wind farm. Moreover, this algorithm is several orders of magnitude faster than the common Kalman filter at a negligible loss in accuracy, with a computational cost of approximately 1.2 s per timestep on an 8-core CPU. Additionally, the estimation solution was tested assuming exclusively readily available measurements in the wind farm. This makes the algorithm feasible for practical implementation. This work provided an essential building block towards a reliable, practicable, closed-loop wind farm control solution with dynamic surrogate models.

6.2. Recommendations

This thesis has made crucial steps towards rendering wind farm control algorithms practicable in commercial wind farms. Nonetheless, many research questions remain unanswered. Several recommendations can be made for future work.

Surrogate models

Firstly, a persistent question in the literature is “how accurate should our surrogate models be for wind farm control?” While surrogate wind farm models now capture the effects of wake steering fairly well, they still often lack in accuracy when modeling wind turbine and wake interactions. Fruitful research directions for surrogate wind farm modeling include:

- A first fruitful research direction for surrogate wind farm models is the inclusion of complex topologies. Surrogate models currently assume a flat ground with a constant surface roughness, which is only somewhat realistic for offshore wind farms. Though, surface elements such as forests and hills have a noticeable effect on the wind distribution and wake development in- and outside of wind farms.
- In close link to the inclusion of complex topologies is the inclusion of non-uniform inflow conditions in surrogate wind farm models. In current surrogate models, the inflow conditions impacting the turbines in a wind farm are assumed to be uniform: each upstream turbine experiences the same inflow wind speed, wind direction and turbulence intensity. In practice, this is far from the truth, especially in large-scale wind farms where two upstream turbines are several kilometers away from one another.

- In view of the work presented by Frederik et al. (2020a), dynamic induction control and helical wake formation by individual pitch control appear as promising technologies to increase the power production of wind farms. However, the optimal excitation signal depends on the inflow conditions and wind farm properties and is not straightforwardly calculated. Therefore, surrogate wind farm models should incorporate the effect of cyclic and individual pitch control on wind turbines and their wakes, such that these models can be leveraged for controller synthesis.
- Currently, wind turbine control algorithms operate the turbine to achieve a delicate balance between the energy yield and structural loading. However, in wind farm control, most concepts only focus on the energy yield, ignoring its impact on structural loads. For practical adoption by the industry, surrogate models should capture the impact of control policies on the structural loads of turbine structures too. This would allow for wind farm control to achieve a similar delicate balance between energy yield and fatigue loads.
- Steady-state surrogate models remain the default choice in wind farm control due to their low complexity and low computational cost. The low computational cost is a prerequisite for real-time controller implementation. To stimulate the usage of higher-fidelity and dynamic models, the increase in computational cost that comes paired with these models must be addressed. This can be done through the usage of efficient numerical solvers, analytical gradients, precalculated solutions and code parallelization, amongst others.
- The inclusion of wake propagation in dynamic surrogate models can greatly increase their flexibility and accuracy. Moreover, the inclusion of wake propagation can prevent long time-averaging or ad-hoc solutions in controller synthesis that would otherwise deal with the large time delays in wind farms. This can also render such models useful for wind farm power production forecasting.
- As seen in Chapter 4, the surrogate wind farm model is most often inaccurate during wind direction changes. The inclusion of flow evolution under such changes in dynamic surrogate models may provide valuable insights in how to minimize performance losses.

Estimation algorithms for wind farm control

Secondly, there is a sparsity in the literature of estimation algorithms for model-based wind farm control. This thesis presents two of such algorithms, one for a steady-state surrogate wind farm model and one for a dynamic surrogate wind farm model. However, many challenges remain concerning these algorithms. Correspondingly, recommendations for future research are:

- Chapter 3 revealed the irregular degree of observability for ambient condition estimation. Therefore, observability analysis should be considered a crucial step in wind farm controller synthesis. In practice, this implies that, during controller synthesis, one has to carefully consider what sensors are available in the wind farm,

what information is contained within them, what variables can be estimated and what variables are desired to be estimated.

- As outlined in Chapter 4, in steady-state-model-based wind farm control, wind turbine measurements are averaged over multiple-minute time windows in attempt to capture the steady-state behavior of the wind farm. However, useful information is lost in this process, which could otherwise be useful in the estimation of the turbulence intensity, for example. Future work should focus on using the information contained in measurements to the full extent, as to maximize the accuracy of the control solution while minimizing the need for additional sensors in the wind farm.
- Chapters 3 and 4 focused on the real-time estimation of the ambient wind speed, wind direction and the wake recovery factor using a steady-state surrogate model and the measurements in the wind farm. An interesting, alternative route is to use these measurements to form local error correction models, as pursued by the group of Prof. Bottasso at TU Munich (e.g., Schreiber et al., 2020). In this framework, the measurements are used to form a *correction model* at each wind turbine to capture dynamics that are not captured in the surrogate model, such as differences in turbine behavior and complex terrain.
- Estimation algorithms employed in closed-loop wind farm control solutions must operate faster than real-time for practical implementation. Therefore, attention much be paid towards the computational tractability of proposed solutions. Inspiration can be gained from the fields of optimization and control for large-scale systems.

Optimization algorithms for wind farm control

In contrast to estimation algorithms, optimization algorithms are receiving an increasing amount of attention in the literature. Though, continued research is much needed on a number of topics, among which are:

- In Chapters 2 and 4 a table of optimal yaw misalignment setpoints as a function of the wind direction, wind speed and wake recovery (or turbulence intensity) is generated. The optimization may lead to strong jumps in the assigned setpoints for small changes in the wind direction or wind speed, leading to excessive yaw motor actuation. Therefore, the setpoints were passed through a Gaussian smoothing function before being applied to the wind farm. However, this post-processing step may lead to setpoints which yield performance losses. Therefore, future research must encompass the question of how to ensure sufficient smoothness in the yaw angle setpoints while guaranteeing optimal performance.
- As wind farms contain an increasing number of wind turbines, the number of degrees of freedom in the optimization grows proportionally. In turn, this leads to an exponential growth in the computational cost involved for the optimization algorithm. While a number of articles address this issue, more research is necessary on

this topic, especially for wind farm control solutions leveraging dynamic surrogate models.

- Distributed wind farm control is a topic that is insufficiently addressed in the literature. Namely, this would not only resolve the issue of the exponential growth in computational cost as the number of turbines inside a wind farm increases, but this would also allow for a more modular approach to wind farm control.

The complete closed-loop control solution

Moreover, two recommendations can be made concerning the complete closed-loop wind farm control solution, namely:

- The closed-loop controller architecture should become the standard for practicable wind farm control algorithms. Namely, surrogate wind farm models are simply not accurate enough to capture all relevant dynamics inside a wind farm. Moreover, the usage of data from downstream turbines' sensors has the potential to maximize the benefits of wind farm control.
- Theoretical notions such as *observability* and *stability* in closed-loop wind farm control solutions, especially those leveraging steady-state surrogate models, are insufficiently understood. Fundamental research on these topics is pivotal to mature the wake steering concept past open-loop control solutions.

Validation of wind farm control technologies

Whereas the previous recommendations focused on controller synthesis, additional recommendations can be made that focus on controller validation. When validating wind farm control algorithms, the following guidelines are suggested:

- High-fidelity simulations subjected to realistic, low-frequency changes in the inflow conditions should become the standard for initial wind farm controller validation. Namely, lower-fidelity simulation models cannot provide sufficient confidence in the efficacy of the tested control solution and high-fidelity simulations with time-invariant inflow conditions insufficiently represent reality.
- Baseline controllers should be picked carefully and validated before using it in comparison with wind farm controllers. Specifically, from experience, the baseline controllers provided in software packages are not necessarily optimal or realistic. In result, the stated benefits of the novel control solution hold little value when compared to this baseline controller.
- Disseminating experimental field data and sharing insights should become the standard for wind farm control trials. This requires academia to incentivise industrial partners through financial and other means (e.g., consultancy, research agreements). This holds more value than the development of more wind farm control solutions through large international research projects.

REFERENCES

- M S Adaramola and P-A Krogstad. Experimental investigation of wake effects on wind turbine performance. *Renewable Energy*, 36(8):2078–2086, 2011. ISSN 0960-1481. URL <http://doi.org/10.1016/j.renene.2011.01.024>.
- J Aho, A Bucksan, J Laks, P A Fleming, Y Jeong, F Dunne, M J Churchfield, L Y Pao, and K E Johnson. A tutorial of wind turbine control for supporting grid frequency through active power control. In *American Control Conference (ACC)*, pages 3120–3131, 2012. URL <http://doi.org/10.1109/ACC.2012.6315180>.
- L E Andersson, B M Doekemeijer, D C van der Hoek, J W van Wingerden, and L Imsland. Adaptation of engineering wake models using Gaussian process regression and high-fidelity simulation data. In *Science of Making Torque from Wind (TORQUE)*, 2020. URL <https://arxiv.org/abs/2003.13323>. In review.
- J R Annoni, P M O Gebraad, A K Scholbrock, P A Fleming, and J W van Wingerden. Analysis of axial-induction-based wind plant control using an engineering and a high-order wind plant model. *Wind Energy*, 19(6):1135–1150, 2016a. ISSN 1099-1824. URL <http://doi.org/10.1002/we.1891>.
- J R Annoni, P M O Gebraad, and P J Seiler. Wind farm flow modeling using an input-output reduced-order model. In *American Control Conference (ACC)*, pages 506–512, 2016b. URL <http://doi.org/10.1109/ACC.2016.7524964>.
- J R Annoni, J Nichols, and P J Seiler. Wind farm modeling and control using dynamic mode decomposition. In *34th Wind Energy Symposium*, AIAA SciTech Forum, pages 1–12. American Institute of Aeronautics and Astronautics (AIAA), 2016c. URL <http://doi.org/10.2514/6.2016-2201>.
- J R Annoni, A K Scholbrock, M J Churchfield, and P A Fleming. Evaluating tilt for wind plants. In *American Control Conference (ACC)*, pages 717–722, 2017. URL <http://doi.org/10.23919/ACC.2017.7963037>.
- J R Annoni, C J Bay, T Taylor, L Y Pao, P A Fleming, and K E Johnson. Efficient optimization of large wind farms for real-time control. In *American Control Conference (ACC)*, pages 6200–6205, 2018a. URL <http://doi.org/10.23919/acc.2018.8430751>.
- J R Annoni, P A Fleming, A K Scholbrock, J Roadman, S Dana, C M Adcock, F Porté-Agel, S Raach, F Haizmann, and D Schlipf. Analysis of control-oriented wake modeling tools using LiDAR field results. *Wind Energy Science*, 3(2):819–831, 2018b. URL <http://doi.org/10.5194/wes-3-819-2018>.
- J R Annoni, C J Bay, K E Johnson, E Dall’Anese, E Quon, T Kemper, and P A Fleming. Wind direction estimation using SCADA data with consensus-based optimization. *Wind Energy Science*, 4(2): 355–368, 2019. URL <http://doi.org/10.5194/wes-4-355-2019>.

- ASTM International. Standard practices for cycle counting in fatigue analysis. Technical Report E1049-85(2017), ASTM International, 2017.
- C Bak, R Bitsche, A Yde, T Kim, M H Hansen, F Zahle, M Gaunaa, J P A A Blasques, M Døssing, J J Wedel Heinen, and T Behrens. Light Rotor: The 10-MW reference wind turbine. In *European Wind Energy Conference and Exhibition (EWEC)*. European Wind Energy Association (EWEA), 2012. URL <https://orbit.dtu.dk/en/publications/light-rotor-the-10-mw-reference-wind-turbine>.
- S Baros and A M Annaswamy. Distributed optimal wind farm control for fatigue load minimization: A consensus approach. *International Journal of Electrical Power & Energy Systems*, 112: 452–459, 2019. ISSN 0142-0615. URL <http://doi.org/10.1016/j.ijepes.2019.04.003>.
- R J Barthelmie, S C Pryor, S T Frandsen, K S Hansen, J G Schepers, K Rados, W Schlez, A Neubert, L E Jensen, and S Neckelmann. Quantifying the impact of wind turbine wakes on power output at offshore wind farms. *Journal of Atmospheric and Oceanic Technology*, 27(8):1302–1317, 2010. URL <https://doi.org/10.1175/2010JTECHA1398.1>.
- J Bartl, F Mühle, and L Sætran. Wind tunnel study on power output and yaw moments for two yaw-controlled model wind turbines. *Wind Energy Science*, 3(2):489–502, 2018. URL <https://doi.org/10.5194/wes-3-489-2018>.
- M Bastankhah and F Porté-Agel. Experimental and theoretical study of wind turbine wakes in yawed conditions. *Journal of Fluid Mechanics*, 806:506—541, 2016. URL <https://doi.org/10.1017/jfm.2016.595>.
- M Bastankhah and F Porté-Agel. Wind farm power optimization via yaw angle control: A wind tunnel study. *Journal of Renewable and Sustainable Energy*, 11(2):023301, 2019. URL <http://doi.org/10.1063/1.5077038>.
- C J Bay, J R Annoni, T Taylor, L Y Pao, and K E Johnson. Active power control for wind farms using distributed model predictive control and nearest neighbor communication. In *American Control Conference (ACC)*, pages 682–687, 2018. URL <http://doi.org/10.23919/acc.2018.8431764>.
- M Bertelè, C L Bottasso, S Cacciola, F Daher Adegas, and S Delport. Wind inflow observation from load harmonics. *Wind Energy Science*, 2(2):615–640, 2017. URL <http://doi.org/10.5194/wes-2-615-2017>.
- L Blomqvist, B W Brook, E C Ellis, P M Kareiva, T Nordhaus, and M Shellenberger. Does the shoe fit? Real versus imagined ecological footprints. *PLOS Biology*, 11(11):1–6, 2013. URL <http://doi.org/10.1371/journal.pbio.1001700>.
- S Boersma, P M O Gebraad, M Vali, B M Doekemeijer, and J W van Wingerden. A control-oriented dynamic wind farm flow model: WFSim. *Journal of Physics: Conference Series*, 753(3):032005, 2016. URL <http://doi.org/10.1088/1742-6596/753/3/032005>.
- S Boersma, B M Doekemeijer, P M O Gebraad, P A Fleming, J R Annoni, A K Scholbrock, J A Fredrik, and J W van Wingerden. A tutorial on control-oriented modeling and control of wind farms. In *American Control Conference (ACC)*, pages 1–18, 2017. URL <http://doi.org/10.23919/ACC.2017.7962923>.

- S Boersma, B M Doekemeijer, M Vali, J Meyers, and J W van Wingerden. A control-oriented dynamic wind farm model: WFSim. *Wind Energy Science*, 3(1):75–95, 2018a. URL <http://doi.org/10.5194/wes-3-75-2018>.
- S Boersma, V Rostampour, B M Doekemeijer, W van Geest, and J W van Wingerden. A constrained model predictive wind farm controller providing active power control: an LES study. *Journal of Physics: Conference Series*, 1037:032023, 2018b. URL <https://doi.org/10.1088/1742-6596/1037/3/032023>.
- S Boersma, V Rostampour, B M Doekemeijer, J W van Wingerden, and T Keviczky. A model predictive wind farm controller with linear parameter-varying models. *IFAC Conference on Nonlinear Model Predictive Control*, 51(20):241–246, 2018c. ISSN 2405-8963. URL <https://doi.org/10.1016/j.ifacol.2018.11.020>.
- S Boersma, B M Doekemeijer, T Keviczky, and J W van Wingerden. Stochastic model predictive control: Uncertainty impact on wind farm power tracking. In *American Control Conference (ACC)*, pages 4167–4172, 2019a. URL <https://doi.org/10.23919/ACC.2019.8814475>.
- S Boersma, B M Doekemeijer, S Siniscalchi-Minna, and J W van Wingerden. A constrained wind farm controller providing secondary frequency regulation and potential power maximization: an LES study. *Renewable Energy*, 134:639–652, 2019b. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2018.11.031>.
- E A Bossanyi. Individual blade pitch control for load reduction. *Wind Energy*, 6(2):119–128, 2003. URL <http://doi.org/10.1002/we.76>.
- E A Bossanyi. CL-Windcon D2.5; chapter 3: Wind farm controller design and testing in LongSim. *CL-Windcon deliverable repository*, 2019. URL <http://www.clwindcon.eu/wp-content/uploads/2020/03/CL-Windcon-D2.5-Integrated-wind-farm-controllers.pdf>. European Horizon 2020 project.
- C L Bottasso and J Schreiber. Online model updating by a wake detector for wind farm control. In *American Control Conference (ACC)*, pages 676–681, 2018. URL <https://doi.org/10.23919/ACC.2018.8431626>.
- F Campagnolo, V Petrović, C L Bottasso, and A Croce. Wind tunnel testing of wake control strategies. In *American Control Conference (ACC)*, pages 513–518, 2016a. URL <https://doi.org/10.1109/ACC.2016.7524965>.
- F Campagnolo, V Petrović, J Schreiber, E M Nanos, A Croce, and C L Bottasso. Wind tunnel testing of a closed-loop wake deflection controller for wind farm power maximization. *Journal of Physics: Conference Series*, 753(3):032006, 2016b. URL <http://doi.org/10.1088/1742-6596/753/3/032006>.
- 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*, 13(14):1–32, 2012a. ISSN 1468-5248. URL <http://doi.org/10.1080/14685248.2012.668191>.
- M J Churchfield, S Lee, P J Moriarty, L A Martínez-Tossas, S Leonardi, G Vijayakumar, and J Brasseur. A large-eddy simulation of wind-plant aerodynamics. In *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2012b. URL <http://doi.org/10.2514/6.2012-537>.

- U Ciri, M A Rotea, and S Leonardi. Model-free control of wind farms: A comparative study between individual and coordinated extremum seeking. *Renewable Energy*, 113:1033–1045, 2017a. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2017.06.065>.
- U Ciri, M A Rotea, C Santoni, and S Leonardi. Large-eddy simulations with extremum-seeking control for individual wind turbine power optimization. *Wind Energy*, 20(9):1617–1634, 2017b. URL <http://doi.org/10.1002/we.2112>.
- U Ciri, S Leonardi, and M A Rotea. Evaluation of log-of-power extremum seeking control for wind turbines using large eddy simulations. *Wind Energy*, 22(7):992–1002, 2019. URL <http://doi.org/10.1002/we.2336>.
- N Cosack. *Fatigue loads monitoring with standard wind turbine signals*. Ph.D. dissertation, University of Stuttgart, 2010. URL <https://d-nb.info/1009926721/34>.
- A Crespo and J Hernández. Turbulence characteristics in wind-turbine wakes. *Journal of Wind Engineering and Industrial Aerodynamics*, 61(1):71–85, 1996. URL [http://doi.org/10.1016/0167-6105\(95\)00033-X](http://doi.org/10.1016/0167-6105(95)00033-X).
- J Å Dahlberg and S E Thor. Power performance and wake effects in the closely spaced Lillgrund offshore wind farm. In *European Offshore Wind Conference and Exhibition (EOW)*. European Wind Energy Association (EWEA), 2009.
- R Damiani, S Dana, J R Annoni, P A Fleming, J Roadman, J van Dam, and K Dykes. Assessment of wind turbine component loads under yaw-offset conditions. *Wind Energy Science*, 3(1):173–189, 2018. URL <http://doi.org/10.5194/wes-3-173-2018>.
- F Díaz-González, M Hau, A Sumper, and O Gomis-Bellmunt. Participation of wind power plants in system frequency control: Review of grid code requirements and control methods. *Renewable and Sustainable Energy Reviews*, 34:551–564, 2014. ISSN 1364-0321. URL <http://doi.org/10.1016/j.rser.2014.03.040>.
- B M Doekemeijer and R Storm. TUDelft FLORIS Github repository, 2019. URL https://github.com/TUDELFT-DataDrivenControl/FLORISSE_M.
- B M Doekemeijer and J W van Wingerden. Observability of the ambient conditions in model-based estimation for wind farm control: A focus on static models. *Wind Energy*, 2020. URL <https://doi.org/10.1002/we.2495>. Early view.
- B M Doekemeijer, J W van Wingerden, S Boersma, and L Y Pao. Enhanced Kalman filtering for a 2D CFD NS wind farm flow model. *Journal of Physics: Conference Series*, 753(5):052015, 2016. URL <https://doi.org/10.1088/1742-6596/753/5/052015>.
- B M Doekemeijer, S Boersma, L Y Pao, and J W van Wingerden. Ensemble Kalman filtering for wind field estimation in wind farms. In *American Control Conference (ACC)*, pages 19–24, 2017. URL <https://doi.org/10.23919/ACC.2017.7962924>.
- B M Doekemeijer, S Boersma, L Y Pao, T Knudsen, and J W van Wingerden. Online model calibration for a simplified LES model in pursuit of real-time closed-loop wind farm control. *Wind Energy Science*, 3(2):749–765, 2018a. URL <https://doi.org/10.5194/wes-3-749-2018>.

- B M Doekemeijer, S Boersma, J W van Wingerden, and L Y Pao. Joint state-parameter estimation for a control-oriented LES wind farm model. *Journal of Physics: Conference Series*, 1037:032013, 2018b. URL <https://doi.org/10.1088/1742-6596/1037/3/032013>.
- B M Doekemeijer, S Boersma, J R Annoni, P A Fleming, and J W van Wingerden. Wind plant controller design. In Paul Veers, editor, *Wind Energy Modeling and Simulation; Volume 1: Atmosphere and plant*, Energy Engineering, chapter 7. The Institution of Engineering and Technology (IET), 2019a. ISBN 978-1-78561-521-4.
- B M Doekemeijer, P A Fleming, and J W van Wingerden. A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model. In *American Control Conference (ACC)*, pages 2825–2836, 2019b. URL <https://doi.org/10.23919/ACC.2019.8815126>.
- B M Doekemeijer, P A Fleming, and J W van Wingerden. Model-based closed-loop wind farm control for power maximization: a large eddy simulation study. In *Conference on Control Technology & Applications*, pages 284–289, 2019c. ISBN 978-1-7281-2767-5. URL <https://doi.org/10.1109/CCTA.2019.8920587>.
- B M Doekemeijer, S Kern, S K Kanev, B Salbert, J Schreiber, F Campagnolo, C L Bottasso, S Maturu, S Schuler, F Wilts, T Neumann, G Potenza, F Calabretta, F Fioretti, and J W van Wingerden. Fields experiments for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science Discussions*, pages 1–22, 2020a.
- B M Doekemeijer, D C van der Hoek, and J W van Wingerden. Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions. *Renewable Energy*, 156:719–730, 2020b. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2020.04.007>.
- B Efron and R J Tibshirani. *An introduction to the bootstrap*. Chapman & Hall, 1993. ISBN 0-412-04231-2.
- E Ela, V Gevorgian, P A Fleming, Y C Zhang, M Singh, E Muljadi, A K Scholbrock, J Aho, A Buckspan, L Y Pao, V Singhvi, A Tuohy, P Pourbeik, D Brooks, and N Bhatt. Active power controls from wind power: bridging the gaps. Technical Report NREL/TP-5D00-60574, National Renewable Energy Laboratory, 2014. URL <https://www.nrel.gov/docs/fy14osti/60574.pdf>.
- Energieleveranciers. Gemiddeld energieverbruik, 2019. URL <https://www.energieleveranciers.nl/gemiddeld-energieverbruik>. Website.
- European Commission. EU reference scenario 2016: Energy, transport and GHG emissions – trends to 2050. Technical report, European Commission, 2016. URL https://ec.europa.eu/energy/sites/ener/files/documents/20160713%20draft_publication_REF2016_v13.pdf.
- European Commission. The EU agricultural outlook for the agricultural markets and income 2017-2030, 2017. URL https://ec.europa.eu/info/sites/info/files/food-farming-fisheries/farming/documents/agricultural-outlook-2017-30_en.pdf. Press release.
- European Commission. The revised Renewable Energy Directive, 2019. URL https://ec.europa.eu/energy/sites/ener/files/documents/directive_renewable_factsheet.pdf. Press release.

- European Commission. Horizon 2020 project repository: Closed loop wind farm control, 2020. URL <https://cordis.europa.eu/project/id/727477>. Accessed April 15, 2020.
- European Environment Agency. EEA policy evaluation: EU law on combustion plants very effective in greening the energy sector, 2019. URL <https://www.eea.europa.eu/highlights/eea-policy-evaluation-eu-law>. News article.
- European Wind Energy Association (EWEA). EWEA fact sheet, 2019. URL <http://www.ewea.org/fileadmin/files/library/publications/statistics/Factsheets.pdf>.
- G Evensen. The Ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367, 2003. ISSN 1616-7228. URL <https://doi.org/10.1007/s10236-003-0036-9>.
- P A Fleming, P M O Gebraad, S Lee, J W van Wingerden, K E Johnson, M J Churchfield, J Michalakes, P Spalart, and P J Moriarty. Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*, 70:211–218, 2014. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2014.02.015>. Special issue on aerodynamics of offshore wind energy systems and wakes.
- P A Fleming, J Aho, P M O Gebraad, L Y Pao, and Y Zhang. Computational fluid dynamics simulation study of active power control in wind plants. In *American Control Conference (ACC)*, pages 1413–1420, 2016. URL <http://doi.org/10.1109/ACC.2016.7525115>.
- P A Fleming, J R Annoni, A K Scholbrock, E Quon, S Dana, S Schreck, S Raach, F Haizmann, and D Schlipf. Full-scale field test of wake steering. *Journal of Physics: Conference Series*, 854(1), 2017a. URL <http://doi.org/10.1088/1742-6596/854/1/012013>.
- P A Fleming, J R Annoni, J J Shah, L Wang, S Ananthan, Z Zhang, K Hutchings, P Wang, W Chen, and L Chen. Field test of wake steering at an offshore wind farm. *Wind Energy Science*, 2(1): 229–239, 2017b. URL <http://doi.org/10.5194/wes-2-229-2017>.
- P A Fleming, J R King, K Dykes, E Simley, J Roadman, A K Scholbrock, P Murphy, J K Lundquist, P J Moriarty, K Fleming, J van Dam, C J Bay, R Mudafort, H Lopez, J Skopek, M Scott, B Ryan, C Guernsey, and D Brake. Initial results from a field campaign of wake steering applied at a commercial wind farm – part 1. *Wind Energy Science*, 4(2):273–285, 2019. URL <http://doi.org/10.5194/wes-4-273-2019>.
- P A Fleming, J R King, E Simley, J Roadman, A Scholbrock, P Murphy, J K Lundquist, P J Moriarty, K Fleming, J van Dam, C J Bay, R Mudafort, D Jager, J Skopek, M Scott, B Ryan, C Guernsey, and D Brake. Continued results from a field campaign of wake steering applied at a commercial wind farm: Part 2. *Wind Energy Science Discussions*, 2020:1–24, 2020. URL <http://doi.org/10.5194/wes-2019-104>.
- J A Frederik, B M Doekemeijer, S P Mulders, and J W van Wingerden. The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms. *Wind Energy*, n/a (n/a), 2020a. URL <http://doi.org/10.1002/we.2513>. Early view.
- J A Frederik, B M Doekemeijer, S P Mulders, and J W van Wingerden. On wind farm wake mixing strategies using dynamic individual pitch control. In *Science of Making Torque from Wind (TORQUE)*, 2020b. In review.

- J A Frederik, R Weber, S Cacciola, F Campagnolo, A Croce, C L Bottasso, and J W van Wingerden. Periodic dynamic induction control of wind farms: proving the potential in simulations and wind tunnel experiments. *Wind Energy Science*, 5:245–257, 2020c. URL <http://doi.org/10.5194/wes-5-245-2020>.
- G Gaspari and S E Cohn. Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757, 1999. ISSN 1477-870X. URL <http://doi.org/10.1002/qj.49712555417>.
- P M O Gebraad and J W van Wingerden. Maximum power-point tracking control for wind farms. *Wind Energy*, 18(3):429–447, 2015. ISSN 1099-1824. URL <http://doi.org/10.1002/we.1706>.
- 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 *American Control Conference (ACC)*, pages 3128–3134, 2014. URL <http://doi.org/10.1109/ACC.2014.6859118>.
- P M O Gebraad, P A Fleming, and J W van Wingerden. Wind turbine wake estimation and control using FLORIDyn, a control-oriented dynamic wind plant model. In *American Control Conference (ACC)*, pages 1702–1708, 2015. URL <http://doi.org/10.1109/ACC.2015.7170978>.
- P M O Gebraad, F W Teeuwisse, J W van Wingerden, P A Fleming, S D Ruben, J R Marden, and L Y Pao. Wind plant power optimization through yaw control using a parametric model for wake effects - a CFD simulation study. *Wind Energy*, 19(1):95–114, 2016. ISSN 1099-1824. URL <http://doi.org/10.1002/we.1822>.
- P M O Gebraad, J J Thomas, A Ning, P A Fleming, and K Dykes. Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control. *Wind Energy*, 20(1):97–107, 2017. ISSN 1099-1824. URL <http://doi.org/10.1002/we.1993>.
- M Giampietro and A Saltelli. Footprints to nowhere. *Ecological Indicators*, 46:610–621, 2014. ISSN 1470-160X. URL <https://doi.org/10.1016/j.ecolind.2014.01.030>.
- S Gillijns, O B Mendoza, J Chandrasekar, B L R De Moor, D S Bernstein, and A Ridley. What is the Ensemble Kalman filter and how well does it work? In *American Control Conference (ACC)*, 2006. URL <http://doi.org/10.1109/ACC.2006.1657419>.
- Global Footprint Network. National footprint and biocapacity accounts, 2019 edition, 2019. URL <https://www.footprintnetwork.org/licenses/public-data-package-free/>. Dataset.
- J P Goit and J Meyers. Optimal control of energy extraction in wind-farm boundary layers. *Journal of Fluid Mechanics*, 768:5–50, 2015. URL <http://doi.org/10.1017/jfm.2015.70>.
- K Hansen. Decision-making based on energy costs: Comparing levelized cost of energy and energy system costs. *Energy Strategy Reviews*, 24:68–82, 2019. ISSN 2211-467X. URL <https://doi.org/10.1016/j.esr.2019.02.003>.
- 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. ISSN 1996-1073. URL <https://doi.org/10.3390/en6020696>.

- L Hickman. Wind myths: Turbines can damage your health, 2012. URL <https://www.theguardian.com/environment/2012/feb/28/wind-turbines-damage-health>. News article.
- P L Houtekamer and H L Mitchell. Ensemble Kalman filtering. *Quarterly Journal of the Royal Meteorological Society*, 131(613):3269–3289, 2005. ISSN 1477-870X. URL <http://doi.org/10.1256/qj.05.135>.
- M F Howland, S K Lele, and J O Dabiri. Wind farm power optimization through wake steering. *Proceedings of National Academy of Sciences*, 116(29):14495–14500, 2019. ISSN 0027-8424. URL <http://doi.org/10.1073/pnas.1903680116>.
- International Energy Agency. World Energy Outlook 2017. Technical report, International Energy Agency, 2017. URL <http://doi.org/10.1787/weo-2017-en>.
- G V Iungo, C Santoni-Ortiz, M Abkar, F Porté-Agel, M A Rotea, and S Leonardi. Data-driven reduced order model for prediction of wind turbine wakes. *Journal of Physics: Conference Series*, 625(1):012009, 2015. URL <http://doi.org/10.1088/1742-6596/625/1/012009>.
- N O Jensen. A note on wind generator interaction. Technical Report RISØ-M-2411, Risø National Laboratory, 1983. URL https://backend.orbit.dtu.dk/ws/portalfiles/portal/55857682/ris_m_2411.pdf.
- Á 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. URL <http://doi.org/10.1002/we.380>.
- R E Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 82(1):35–45, 1960. URL <https://www.cs.unc.edu/~welch/kalman/media/pdf/Kalman1960.pdf>.
- S K Kanev. Dynamic wake steering and its impact on wind farm power production and yaw actuator duty. *Renewable Energy*, 146:9–15, 2020. ISSN 0960-1481. URL <http://doi.org/10.1016/j.renene.2019.06.122>.
- S K Kanev, F J Savenije, and W P Engels. Active wake control: An approach to optimize the lifetime operation of wind farms. *Wind Energy*, 21(7):488–501, 2018. URL <http://doi.org/10.1002/we.2173>.
- I Katic, J Højstrup, and N O Jensen. *A simple model for cluster efficiency*, volume 1, pages 407–410. European Wind Energy Association (EWEA), 1987. URL <https://orbit.dtu.dk/en/publications/a-simple-model-for-cluster-efficiency>.
- J Kazda and N A Cutulus. Model-optimized dispatch for closed-loop power control of waked wind farms. *IEEE Transactions on Control Systems Technology*, pages 1–8, 2019. ISSN 1558-0865. URL <http://doi.org/10.1109/TCST.2019.2923779>.
- A C Kheirabadi and R Nagamune. A quantitative review of wind farm control with the objective of wind farm power maximization. *Journal of Wind Engineering and Industrial Aerodynamics*, 192:45–73, 2019. ISSN 0167-6105. URL <https://doi.org/10.1016/j.jweia.2019.06.015>.

- T Knudsen, T Bak, and M Svenstrup. Survey of wind farm control – power and fatigue optimization. *Wind Energy*, 18(8):1333–1351, 2015. ISSN 1099-1824. URL <http://doi.org/10.1002/we.1760>.
- K A Kragh and P A Fleming. Rotor speed dependent yaw control of wind turbines based on empirical data. In *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, pages 1–9, 2012. URL <http://doi.org/10.2514/6.2012-1018>.
- 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*, pages 1207–1214, 2013. ISSN 1558-0865. URL <https://doi.org/10.1109/TCST.2013.2257780>.
- L A Martínez-Tossas, J R Annoni, P A Fleming, and M J Churchfield. The aerodynamics of the curled wake: A simplified model in view of flow control. *Wind Energy Science*, 4(1):127–138, 2019. URL <http://doi.org/10.5194/wes-4-127-2019>.
- V Masson-Delmotte, P Zhai, and H O Pörtner and et al. Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Technical report, Intergovernmental Panel on Climate Change (IPCC), 2019. URL https://www.ipcc.ch/site/assets/uploads/sites/2/2019/06/SR15_Full_Report_High_Res.pdf.
- N Mittelmeier, T Blodau, and M Kühn. Monitoring offshore wind farm power performance with SCADA data and an advanced wake model. *Wind Energy Science*, 2(1):175–187, 2017. URL <http://doi.org/10.5194/wes-2-175-2017>.
- S P Mulders. *Wind turbine control: advances for load mitigations and hydraulic drivetrains*. Ph.D. dissertation, Delft University of Technology, 2020.
- S P Mulders and J W van Wingerden. On the importance of the azimuth offset in a combined 1P and 2P SISO IPC implementation for wind turbine fatigue load reductions. In *American Control Conference (ACC)*, pages 3506–3511, 2019. URL <http://doi.org/10.23919/ACC.2019.8814829>.
- W Munters and J Meyers. An optimal control framework for dynamic induction control of wind farms and their interaction with the atmospheric boundary layer. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 375(2091), 2017. ISSN 1364-503X. URL <http://doi.org/10.1098/rsta.2016.0100>.
- W Munters and J Meyers. Towards practical dynamic induction control of wind farms: analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines. *Wind Energy Science*, 3(1):409–425, 2018. URL <http://doi.org/10.5194/wes-3-409-2018>.
- National Renewable Energy Laboratory (NREL). NWTC Information Portal (OpenFAST), 2019. URL <https://nwtc.nrel.gov/OpenFAST>.
- T Obdam, L Rademakers, H Braam, and P Eecen. Estimating costs of operation & maintenance for offshore wind farms. Technical Report ECN-M-07-120, Energieonderzoek Centrum Nederland, 2007. URL <https://publicaties.ecn.nl/PdfFetch.aspx?nr=ECN-M--07-120>.

- R Ortega, F Mancilla-David, and F Jaramillo. A globally convergent wind speed estimator for wind turbine systems. *International Journal of Adaptive Control and Signal Processing*, 27(5):413–425, 2013. URL <http://doi.org/10.1002/acs.2319>.
- R J Pachauri, M R Allen, and V Barros et al. Climate Change 2014: Synthesis Report. Contribution of working groups I, II and III to the fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, Intergovernmental Panel on Climate Change (IPCC), 2014. URL https://www.ipcc.ch/site/assets/uploads/2018/02/SYR_AR5_FINAL_full.pdf.
- L Y Pao and K E Johnson. A tutorial on the dynamics and control of wind turbines and wind farms. In *American Control Conference (ACC)*, pages 2076–2089, 2009. URL <http://doi.org/10.1109/ACC.2009.5160195>.
- J Park, S D Kwon, and K H Law. A data-driven approach for cooperative wind farm control. In *American Control Conference (ACC)*, pages 525–530, 2016. URL <http://doi.org/10.1109/ACC.2016.7524967>.
- R E Petrie. *Localization in the Ensemble Kalman filter*. Master's dissertation, University of Reading, 2008. URL <http://www.met.rdg.ac.uk/mscdissertations/LocalizationintheensembleKalmanFilter.pdf>.
- A Philbin and W Raillant-Clark. ICAO council adopts new CO2 emissions standard for aircraft, 2017. URL <https://www.icao.int/Newsroom/Pages/ICAO-Council-adopts-new-CO2-emissions-standard-for-aircraft.aspx>. News article.
- S Raach, S Boersma, B M Doekemeijer, J W van Wingerden, and P W Cheng. LiDAR-based closed-loop wake redirection in high-fidelity simulation. *Journal of Physics: Conference Series*, 1037:032016, 2018. URL <http://doi.org/10.1088/1742-6596/1037/3/032016>.
- S Raach, B M Doekemeijer, S Boersma, J W van Wingerden, and P W Cheng. Feedforward-feedback wake redirection for wind farm control. *Wind Energy Science Discussions*, pages 1–18, 2019. URL <https://doi.org/10.5194/wes-2019-54>.
- H Mendez Reyes, S K Kanev, B M Doekemeijer, and J W van Wingerden. Validation of a lookup-table approach to modeling turbine fatigue loads in wind farms under active wake control. *Wind Energy Science*, 4(4):549–561, 2019. URL <http://doi.org/10.5194/wes-4-549-2019>.
- Rijksdienst voor Ondernemend Nederland. Stimulering duurzame energieproductie: Windenergie op zee, 2019. URL <https://www.rvo.nl/subsidies-regelingen/stimulering-duurzame-energieproductie/categorie%C3%ABn/windenergie-op-zee>. Website.
- Rijksoverheid. Klimaatakkoord. Technical report, Rijksoverheid, 2019. URL <https://www.rijksoverheid.nl/binaries/rijksoverheid/documenten/rapporten/2019/06/28/klimaatakkoord/klimaatakkoord.pdf>.
- B Ritter, E Mora, A Schild, B M Doekemeijer, and U Konigorski. Adaptive master-slave Cubature Kalman filters subject to state inequality constraints for wind turbine state estimation. In *American Control Conference (ACC)*, pages 3482–3487, 2019. URL <https://doi.org/10.23919/ACC.2019.8814411>.

- M A Rotea. Dynamic programming framework for wind power maximization. *IFAC World Congress*, 47(3):3639–3644, 2014. ISSN 1474-6670. URL <http://doi.org/10.3182/20140824-6-ZA-1003.02071>.
- A Rott, B M Doekemeijer, J Seifert, J W van Wingerden, and M Kühn. Robust active wake control in consideration of wind direction variability and uncertainty. *Wind Energy Science*, 3(2):869–882, 2018. URL <http://doi.org/10.5194/wes-3-869-2018>.
- S Kern et al. CL-Windcon D3.2: definition of field-testing conditions. *CL-Windcon deliverable repository*, 2017. URL <http://www.clwindcon.eu/wp-content/uploads/2017/03/CL-Windcon-D3.2-Definition-field-testing.pdf>. European Horizon 2020 project.
- C Santoni, U Ciri, M A Rotea, and S Leonardi. Development of a high fidelity CFD code for wind farm control. In *American Control Conference (ACC)*, pages 1715–1720, 2015. URL <http://doi.org/10.1109/ACC.2015.7170980>.
- J Schottler, A Hölling, J Peinke, and M Hölling. Wind tunnel tests on controllable model wind turbines in yaw. In *34th Wind Energy Symposium*, AIAA SciTech Forum, pages 1–8. American Institute of Aeronautics and Astronautics (AIAA), 2016. URL <https://doi.org/10.2514/6.2016-1523>.
- J Schreiber, E M Nanos, F Campagnolo, and C L Bottasso. Verification and calibration of a reduced order wind farm model by wind tunnel experiments. *Journal of Physics: Conference Series*, 854(1):012041, 2017. URL <https://doi.org/10.1088/1742-6596/854/1/012041>.
- J Schreiber, C L Bottasso, B, Salbert, and F Campagnolo. Improving wind farm flow models by learning from operational data. *Wind Energy Science*, 5(2):647–673, 2020. URL <https://doi.org/10.5194/wes-5-647-2020>.
- C R Shapiro, P Bauweraerts, J Meyers, C Meneveau, and D F Gayme. Model-based receding horizon control of wind farms for secondary frequency regulation. *Wind Energy*, 20(7):1261–1275, 2017a. ISSN 1099-1824. URL <http://doi.org/10.1002/we.2093>.
- C R Shapiro, J Meyers, C Meneveau, and D F Gayme. Dynamic wake modeling and state estimation for improved model-based receding horizon control of wind farms. In *American Control Conference (ACC)*, pages 709–716, 2017b. URL <https://doi.org/10.23919/ACC.2017.7963036>.
- C R Shapiro, J Meyers, C Meneveau, and D F Gayme. Wind farms providing secondary frequency regulation: evaluating the performance of model-based receding horizon control. *Wind Energy Science*, 3(1):11–24, 2018. URL <https://doi.org/10.5194/wes-3-11-2018>.
- Siemens-Gamesa Renewable Energy. Wake Adapt solution, 2019. URL <https://www.youtube.com/watch?v=UP2yJqGQ2FU>. Promotional video; product advertisement.
- E Simley and L Y Pao. Evaluation of a wind speed estimator for effective hub-height and shear components. *Wind Energy*, 19(1):167–184, 2016. URL <http://doi.org/10.1002/we.1817>.
- E Simley, P A Fleming, and J R King. Design and analysis of a wake steering controller with wind direction variability. *Wind Energy Science*, 5(2):451–468, 2020. URL <https://doi.org/10.5194/wes-5-451-2020>.

- S Siniscalchi-Minna, F D Bianchi, and C Ocampo-Martinez. Predictive control of wind farms based on lexicographic minimizers for power reserve maximization. In *American Control Conference (ACC)*, pages 701–706, 2018. URL <https://doi.org/10.23919/ACC.2018.8431865>.
- S Siniscalchi-Minna, F D Bianchi, M De-Prada-Gil, and C Ocampo-Martinez. A wind farm control strategy for power reserve maximization. *Renewable Energy*, 131:37–44, 2019. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2018.06.112>.
- 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. URL <https://doi.org/10.1002/we.1594>.
- M N Soltani, T Knudsen, M Svenstrup, R Wisniewski, P Brath, R Ortega, and K E Johnson. Estimation of rotor effective wind speed: A comparison. *IEEE Transactions on Control Systems Technology*, 21(4):1155–1167, 2013. ISSN 1063-6536. URL <https://doi.org/10.1109/TCST.2013.2260751>.
- J N Sørensen and W Z Shen. Numerical modeling of wind turbine wakes. *Journal of Fluid Mechanics*, 124(2):393–399, 2002. URL <https://doi.org/10.1115/1.1471361>.
- P E Sørensen, A D Hansen, K Thomsen, T Buhl, P E Morthorst, L H Nielsen, F Iov, F Blaabjerg, H A Nielsen, H Madsen, and M H Donovan. Operation and control of large wind turbines and wind farms. Technical Report 1532(EN), Forskningscenter Risø, 2005. URL <https://orbit.dtu.dk/en/publications/operation-and-control-of-large-wind-turbines-and-wind-farms-final>.
- T Sørensen, P Nielsen, and M L Thøgersen. Recalibrating wind turbine wake model parameters – validating the wake model performance for large offshore wind farms. In *European Wind Energy Conference and Exhibition (EWEC)*. European Wind Energy Association (EWEA), 2006. URL <https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.610.5404>.
- V Spudić, C Conte, M Baotić, and M Morari. Cooperative distributed model predictive control for wind farms. *Optimal Control Applications and Methods*, 36(3):333–352, 2015. URL <https://doi.org/10.1002/oca.2136>.
- N Troldborg. *Actuator line modeling of wind turbine wakes*. Ph.D. dissertation, Technical University of Denmark, 2008.
- United Nations. Adoption of the Paris Agreement. Technical Report FCCC/CP/2015/L.9/Rev.1, United Nations, 2016. URL <http://unfccc.int/resource/docs/2015/cop21/eng/109r01.pdf>.
- M Vali, V Petrović, S Boersma, J W van Wingerden, and M Kühn. Adjoint-based model predictive control of wind farms: Beyond the quasi steady-state power maximization. *IFAC World Congress*, 50(1):4510–4515, 2017. ISSN 2405-8963. URL <https://doi.org/10.1016/j.ifacol.2017.08.382>.
- D C van der Hoek, B M Doekemeijer, L E Andersson, and J W van Wingerden. Predicting the effect of wake steering on the annual energy production of a large wind farm using high-fidelity simulations and Gaussian process regression. In *Science of Making Torque from Wind (TORQUE)*, 2020. In review.

- M T van Dijk, J W van Wingerden, T Ashuri, and Y Li. Wind farm multi-objective wake redirection for optimizing power production and loads. *Energy*, 121:561–569, 2017. ISSN 0360-5442. URL <https://doi.org/10.1016/j.energy.2017.01.051>.
- G A M van Kuik. The Lanchester–Betz–Joukowski limit. *Wind Energy*, 10(3):289–291, 2007. URL <http://doi.org/10.1002/we.218>.
- J W van Wingerden, L Y Pao, J Aho, and P A Fleming. Active power control of waked wind farms. In *IFAC World Congress*, volume 50, pages 4484–4491, 2017. URL <https://doi.org/10.1016/j.ifacol.2017.08.378>.
- J W van Wingerden, P A Fleming, T Göçmen, I Eguinoa, B M Doekemeijer, K Dykes, M Lawson, E Simley, J R King, D Astrain, M Iribas, C L Bottasso, J Meyers, S Raach, K Kölle, and G Giebel. Expert elicitation on wind farm control. In *Science of Making Torque from Wind (TORQUE)*, 2020. In review.
- L Vollmer. *Influence of atmospheric stability on wind farm control*. PhD thesis, University of Oldenburg, 2018.
- J W Wagenaar, L A H Machielse, and J G Schepers. Controlling wind in ecn's scaled wind farm. In *European Wind Energy Conference and Exhibition (EWECE)*. European Wind Energy Association (EWEA), 2012. URL <https://publications.tno.nl/publication/34631412/306uxY/m12007.pdf>.
- E A Wan and R Van Der Merwe. The Unscented Kalman filter for nonlinear estimation. In *IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, pages 153–158, 2000. URL <https://doi.org/10.1109/ASSPCC.2000.882463>.
- J Yeung. Rising sea levels threaten hundreds of millions – and it's much worse than we thought, 2019. URL <https://edition.cnn.com/2019/10/30/world/rising-sea-cities-study-intl-hnk-scli-sci/index.html>. News article.

LIST OF ABBREVIATIONS

2D	Two-dimensional
ABL	Atmospheric Boundary Layer
ADM	(Generalized) Actuator Disk Model
AEP	Annual Energy Production
ALM	Actuator Line Model
APC	Active Power Control
BEM	Blade Element Momentum
CL	Closed-Loop
CL-Windcon	Closed-Loop Wind Farm Control
Cov	Covariance
CPU	Central Processing Unit
CU Boulder	University of Colorado Boulder
DEL	Damage Equivalent Load
DTU	Technical University of Denmark
EGP	ENEL Green Power
EnKF	Ensemble Kalman Filter
EWEA	European Wind Energy Association
ExKF	Extended Kalman Filter
FLORIS	Flow Redirection and Induction in Steady-State
GE	General Electric
IPCC	International Panel on Climate Change
KF	Kalman Filter
LCOE	Levelized Cost Of Energy
LiDAR	Light Detection And Ranging
LES	Large-Eddy Simulation
NREL	U.S. National Renewable Energy Laboratory
NS	Navier-Stokes
OL	Open-Loop
SCADA	Supervisory Control And Data Acquisition
SOWFA	Simulator for Onshore/Offshore Wind Farm Applications
TI	Turbulence Intensity
TUD, TUDelft	Delft University of Technology
UKF	Unscented Kalman Filter
WTG	Wind Turbine Generator
WD	Wind Direction
WF	Wind Farm
WFC	Wind Farm Control
WFSim	Wind Farm Simulator
WS	Wind Speed

A

APPENDIX: ADDITIONAL FIELD EXPERIMENT LOOK-UP TABLES

This appendix presents additional yaw setpoint look-up tables for the wake steering field experiment described in Chapter 2. The turbine yaw setpoints were optimized for a large range of inflow conditions as described in Section 2.3.2. Figure 2.6 previously showed the optimal yaw setpoints for a low turbulence intensity of 7.5%. This appendix shows the optimal yaw setpoints for turbulence intensities of 13.5% and 18.0%.

The optimal turbine yaw setpoints for a turbulence intensity of 13.5% are shown in Figure A.1. Compared to the situation with a turbulence intensity of 7.5%, the forecasted performance gains notably reduce. A higher ambient turbulence leads to more wake recovery and thus the benefits of wake steering become less apparent. The optimal turbine yaw setpoints for a turbulence intensity of 18.0% are shown in Figure A.2. Compared to the situations with turbulence intensities of 7.5% and 13.5%, the gains are very small. In practice, these gains are expected to drown in statistical uncertainty.

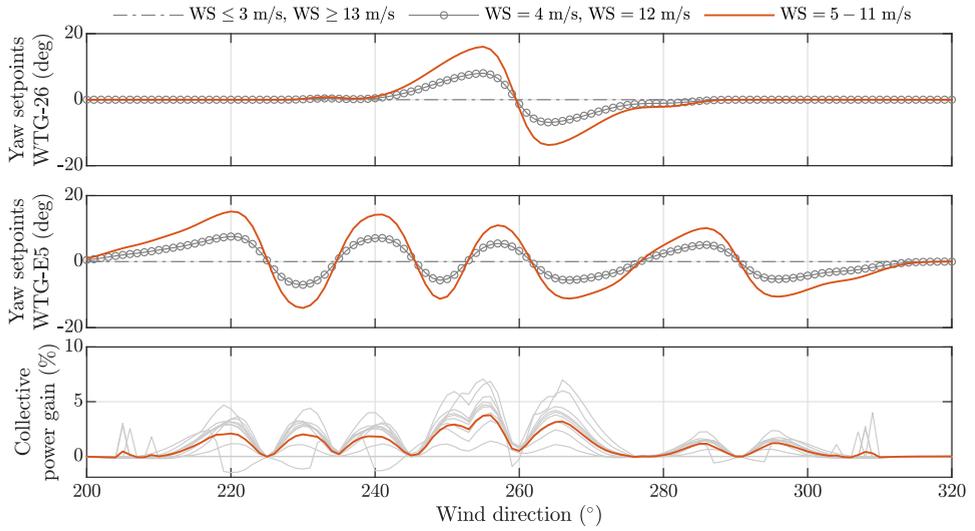


Figure A.1: The optimal turbine yaw angle setpoints for WTG 26 and E5 for a freestream turbulence intensity of 13.5%. The collective power gain of WTG 26, WTG E5 and the downstream machine (WTG 10, 11, 12, or 31) is shown in the bottom plot.

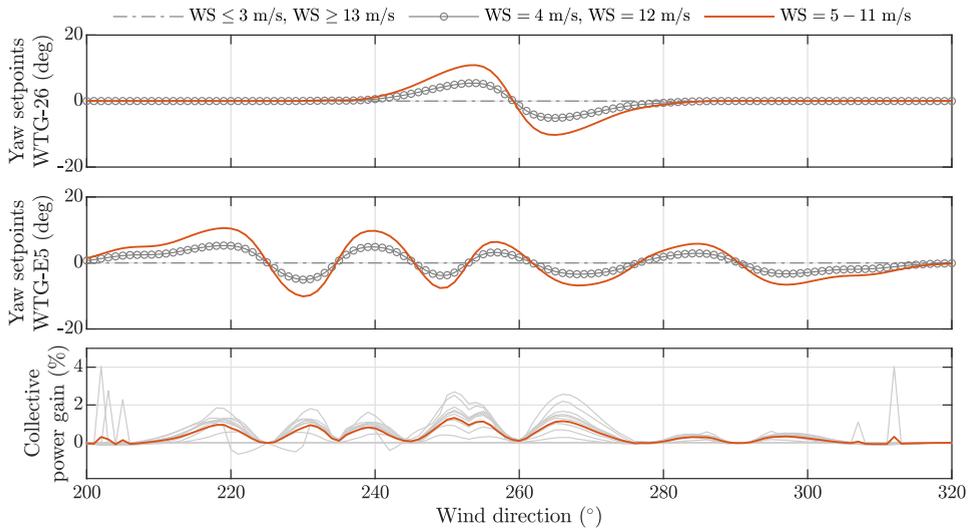


Figure A.2: The optimal turbine yaw angle setpoints for WTG 26 and E5 for a freestream turbulence intensity of 18.0%. The collective power gain of WTG 26, WTG E5 and the downstream machine (WTG 10, 11, 12, or 31) is shown in the bottom plot.

B

APPENDIX: YAW-POWER RELATIONSHIP FOR A GE 1.5S TURBINE

This appendix presents additional information on the turbine yaw-power behavior for the field experiment described in Chapter 2. The experimental results from Section 2.5 indicate that negative yaw misalignment in WTG 26 leads to very small losses and sometimes even to a power gain compared to aligned operation. This behavior is verified by studying experimental data from a different GE 1.5s turbine inside the Sedini wind farm that is not included in the wake steering experiments: WTG 30. SCADA data of this turbine is used to plot the normalized power production of the turbine against its yaw misalignment angle, shown in Figure B.1. This figure shows that there is practically no decrease in power production when misaligning the turbine in the negative direction by less than 10° . This is in agreement with the behavior seen in WTG 26 and explains the large gains around the $260 - 280^\circ$ region in the field experiments shown in Figure 2.11: wake steering effectively comes “for free” here.

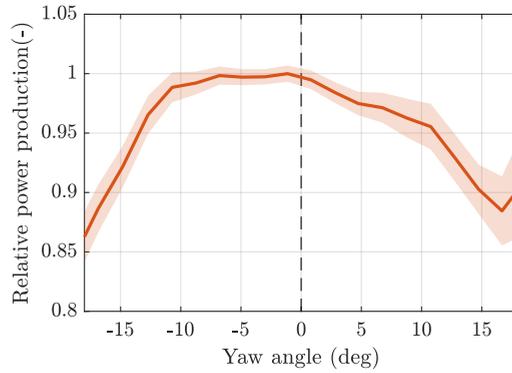


Figure B.1: Relationship between the normalized power production and the yaw misalignment angle for an arbitrary GE 1.5s wind turbine in the Sedini wind farm. The data was collected for the range of 6 m/s to 12 m/s wind speeds. The asymmetry is clearly seen. Moreover, negative yaw misalignment shows a much smaller loss or even a very slight gain in power production compared to positive yaw misalignment.

C

APPENDIX: ADDRESSING IRREGULARITIES IN THE DEGREE OF OBSERVABILITY

This appendix provides a brief discussion on the issues that may arise in the measure of observability defined in Chapter 3. A solution is also presented.

Equation 3.10 provides a clear measure for the degree of observability of a particular situation. With this formulation, evaluated ambient conditions far away from the true ambient conditions (e.g., $\|\Delta\phi\| \gg 0$) that yield a low estimation error J are penalized heavily. Namely, the nominator is small and the denominator is large, leading to a low value of \mathcal{O} . In such a situation, it is unclear what the true ambient conditions are based on the measurements available. These situations result in a low degree of observability. Alternatively, situations with a high cost far away from the true ambient conditions result in a high degree of observability.

However, by simply dividing the cost function J over the distance between the evaluated and true ambient conditions leads to undesired behaviour near the true ambient conditions (e.g., $\Delta\phi \approx 0$). For example, a singularity arises when the evaluated ambient conditions $\hat{\phi}$, \hat{U}_∞ and \hat{I}_∞ are exactly the true ambient conditions ϕ , U_∞ and I_∞ , respectively. Namely, then

$$\mathcal{M}(\phi, U_\infty, I_\infty) = \frac{0}{0} = \text{undefined.}$$

Similarly, when the evaluated conditions are very close to the true conditions, it becomes difficult to envision what the function of \mathcal{M} will look like. For example, if $J = 0$ at $\Delta\hat{\phi} = 0.2^\circ$, then the situation would turn out to be unobservable. This is because one cannot distinguish the true ambient condition ($\phi = 0^\circ$) from a different evaluated condition ($\hat{\phi} = 0.2^\circ$). Clearly, this should not yield an unobservable situation and a situation

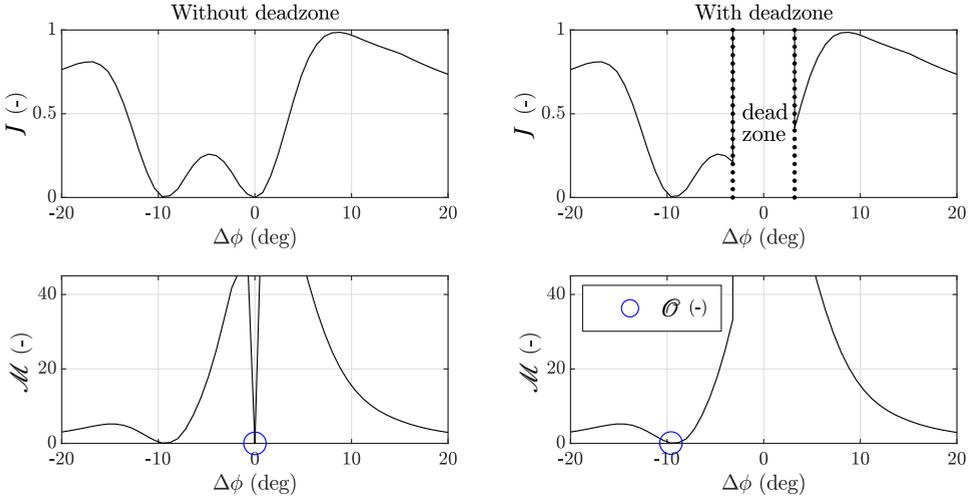


Figure C.1: This figure depicts the issue when $\Delta\phi \approx 0$ for the calculation of \mathcal{M} using Equation 3.10. The cost function shown here refers back to the 2-turbine wind farm previously discussed in Section 3.3.1 with $\lambda_\phi = \lambda_U = 0$. In the top-left figure, J , the mean-squared-error in turbine power signals, is plotted as a function of the hypothesized wind direction. When looking at the left two subplots (without deadzone), the observability is found to be $\Theta = 0$ at $\Delta\phi = 0^\circ$ due to a singularity in the equation. Naturally, the error at $\Delta\phi = 0^\circ$ should always be zero by definition of J , and this in no scenario should be defined as unobservability. Therefore, with the introduction of the deadzone (right two subplots), this issue is resolved. Basically, the deadzone enforces that the situation will never be unobservable if $J \approx 0$ in close proximity of the true solution.

where J is very low “close enough” to the true conditions should not negatively impact the observability of the situation. To address this issue, a “deadzone” is introduced for \mathcal{M} in proximity of the true ambient conditions. This deadzone enforces observability when the evaluated ambient conditions are “close enough” to the true ambient conditions. This can be seen as the upper formula in Equation 3.10, in which $\mathcal{M} = \infty$ within the deadzone region. The effect of a deadzone is visualized in Figure C.1. This deadzone resolves the issues related to singularities and numerical sensitivities.

CURRICULUM VITAE

Born

1993 Born in Dordrecht, The Netherlands.

Education

2005–2011 Highschool
Insula College Thuredrecht, Dordrecht (2005-2011)

2011–2014 Bachelor's degree in in Aerospace Engineering (*cum laude*)
Delft University of Technology, The Netherlands

2014–2016 Master's degree in Systems & Control Engineering (*cum laude*)
Delft University of Technology, The Netherlands

2016–2020 Doctoral degree in Control Engineering
Delft University of Technology, The Netherlands
Thesis title: Closing the loop in model-based wind farm control
Promotors: Prof. dr. ir. J.W. van Wingerden, Prof. dr. ir. M. Verhaegen

Honors & Awards

2016 PhD Best Poster Award (1st place), TUDelft Energy Initiative: Poster Event Delft, The Netherlands

2018 Invited Speaker of prof. L. Y. Pao at CU Boulder, Seminar on the topics of Robotics, Control and Dynamical Systems (RCDS) Boulder, USA

2019 Outstanding Student Paper Award, IEE Conference on Control and Applications 2019 Hong Kong, China

2019 Invited speaker at the 5th Wind Energy Systems Engineering Workshop Pamplona, Spain

2018-2020 Co- and lead organizer roles for invited sessions on wind energy control, American Conference 2018, 2019 and 2020 Milwaukee, USA (2018); Philadelphia, USA (2019); Denver, USA (2020)

LIST OF PUBLICATIONS

Book chapters

1. * B M Doekemeijer, S Boersma, J R Annoni, P A Fleming, and J W van Wingerden. Wind plant controller design. In Paul Veers, editor, *Wind Energy Modeling and Simulation; Volume 1: Atmosphere and plant*, Energy Engineering, chapter 7. The Institution of Engineering and Technology (IET), 2019a. ISBN 978-1-78561-521-4

Journal publications

10. * B M Doekemeijer, S Kern, S K Kanev, B Salbert, J Schreiber, F Campagnolo, C L Bottasso, S Maturu, S Schuler, F Wilts, T Neumann, G Potenza, F Calabretta, F Fioretti, and J W van Wingerden. Fields experiments for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science Discussions*, pages 1–22, 2020a
9. * B M Doekemeijer, D C van der Hoek, and J W van Wingerden. Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions. *Renewable Energy*, 156:719–730, 2020b. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2020.04.007>
8. * B M Doekemeijer and J W van Wingerden. Observability of the ambient conditions in model-based estimation for wind farm control: A focus on static models. *Wind Energy*, 2020. URL <https://doi.org/10.1002/we.2495>. Early view
7. * B M Doekemeijer, S Boersma, L Y Pao, T Knudsen, and J W van Wingerden. Online model calibration for a simplified LES model in pursuit of real-time closed-loop wind farm control. *Wind Energy Science*, 3(2):749–765, 2018a. URL <https://doi.org/10.5194/wes-3-749-2018>
6. J A Frederik, B M Doekemeijer, S P Mulders, and J W van Wingerden. The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms. *Wind Energy*, n/a(n/a), 2020a. URL <http://doi.org/10.1002/we.2513>. Early view
5. H Mendez Reyes, S K Kanev, B M Doekemeijer, and J W van Wingerden. Validation of a lookup-table approach to modeling turbine fatigue loads in wind farms under active wake control. *Wind Energy Science*, 4(4):549–561, 2019. URL <http://doi.org/10.5194/wes-4-549-2019>
4. S Raach, B M Doekemeijer, S Boersma, J W van Wingerden, and P W Cheng. Feedforward-feedback wake redirection for wind farm control. *Wind Energy Science Discussions*, pages 1–18, 2019. URL <https://doi.org/10.5194/wes-2019-54>
3. S Boersma, B M Doekemeijer, S Siniscalchi-Minna, and J W van Wingerden. A constrained wind farm controller providing secondary frequency regulation and potential power maximization: an LES study. *Renewable Energy*, 134:639–652, 2019b. ISSN 0960-1481. URL <https://doi.org/10.1016/j.renene.2018.11.031>
2. A Rott, B M Doekemeijer, J Seifert, J W van Wingerden, and M Kühn. Robust active wake control in consideration of wind direction variability and uncertainty. *Wind Energy Science*, 3(2):869–882, 2018. URL <http://doi.org/10.5194/wes-3-869-2018>

LIST OF PUBLICATIONS

1. S Boersma, B M Doekemeijer, M Vali, J Meyers, and J W van Wingerden. A control-oriented dynamic wind farm model: WFSim. *Wind Energy Science*, 3(1):75–95, 2018a. URL <http://doi.org/10.5194/wes-3-75-2018>

Tutorial publications

3. * B M Doekemeijer, P A Fleming, and J W van Wingerden. A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model. In *American Control Conference (ACC)*, pages 2825–2836, 2019b. URL <https://doi.org/10.23919/ACC.2019.8815126>
2. J W van Wingerden, P A Fleming, T Göçmen, I Eguinoa, B M Doekemeijer, K Dykes, M Lawson, E Simley, J R King, D Astrain, M Iribas, C L Bottasso, J Meyers, S Raach, K Kölle, and G Giebel. Expert elicitation on wind farm control. In *Science of Making Torque from Wind (TORQUE)*, 2020. In review
1. S Boersma, B M Doekemeijer, P M O Gebraad, P A Fleming, J R Annoni, A K Scholbrock, J A Frederik, and J W van Wingerden. A tutorial on control-oriented modeling and control of wind farms. In *American Control Conference (ACC)*, pages 1–18, 2017. URL <http://doi.org/10.23919/ACC.2017.7962923>

Conference publications

13. * B M Doekemeijer, P A Fleming, and J W van Wingerden. Model-based closed-loop wind farm control for power maximization: a large eddy simulation study. In *Conference on Control Technology & Applications*, pages 284–289, 2019c. ISBN 978-1-7281-2767-5. URL <https://doi.org/10.1109/CCTA.2019.8920587>
12. * B M Doekemeijer, S Boersma, J W van Wingerden, and L Y Pao. Joint state-parameter estimation for a control-oriented LES wind farm model. *Journal of Physics: Conference Series*, 1037:032013, 2018b. URL <https://doi.org/10.1088/1742-6596/1037/3/032013>
11. * B M Doekemeijer, S Boersma, L Y Pao, and J W van Wingerden. Ensemble Kalman filtering for wind field estimation in wind farms. In *American Control Conference (ACC)*, pages 19–24, 2017. URL <https://doi.org/10.23919/ACC.2017.7962924>
10. * B M Doekemeijer, J W van Wingerden, S Boersma, and L Y Pao. Enhanced Kalman filtering for a 2D CFD NS wind farm flow model. *Journal of Physics: Conference Series*, 753(5):052015, 2016. URL <https://doi.org/10.1088/1742-6596/753/5/052015>
9. J A Frederik, B M Doekemeijer, S P Mulders, and J W van Wingerden. On wind farm wake mixing strategies using dynamic individual pitch control. In *Science of Making Torque from Wind (TORQUE)*, 2020b. In review
8. D C van der Hoek, B M Doekemeijer, L E Andersson, and J W van Wingerden. Predicting the effect of wake steering on the annual energy production of a large wind farm using high-fidelity simulations and Gaussian process regression. In *Science of Making Torque from Wind (TORQUE)*, 2020. In review
7. L E Andersson, B M Doekemeijer, D C van der Hoek, J W van Wingerden, and L Imsland. Adaptation of engineering wake models using Gaussian process regression and high-fidelity simulation data. In *Science of Making Torque from Wind (TORQUE)*, 2020. URL <https://arxiv.org/abs/2003.13323>. In review

6. S Boersma, B M Doekemeijer, T Keviczky, and J W van Wingerden. Stochastic model predictive control: Uncertainty impact on wind farm power tracking. In *American Control Conference (ACC)*, pages 4167–4172, 2019a. URL <https://doi.org/10.23919/ACC.2019.8814475>
5. B Ritter, E Mora, A Schild, B M Doekemeijer, and U Konigorski. Adaptive master-slave Cubature Kalman filters subject to state inequality constraints for wind turbine state estimation. In *American Control Conference (ACC)*, pages 3482–3487, 2019. URL <https://doi.org/10.23919/ACC.2019.8814411>
4. S Boersma, V Rostampour, B M Doekemeijer, W van Geest, and J W van Wingerden. A constrained model predictive wind farm controller providing active power control: an LES study. *Journal of Physics: Conference Series*, 1037:032023, 2018b. URL <https://doi.org/10.1088/1742-6596/1037/3/032023>
3. S Boersma, V Rostampour, B M Doekemeijer, J W van Wingerden, and T Keviczky. A model predictive wind farm controller with linear parameter-varying models. *IFAC Conference on Nonlinear Model Predictive Control*, 51(20):241–246, 2018c. ISSN 2405-8963. URL <https://doi.org/10.1016/j.ifacol.2018.11.020>
2. S Raach, S Boersma, B M Doekemeijer, J W van Wingerden, and P W Cheng. LiDAR-based closed-loop wake redirection in high-fidelity simulation. *Journal of Physics: Conference Series*, 1037:032016, 2018. URL <http://doi.org/10.1088/1742-6596/1037/3/032016>
1. S Boersma, P M O Gebraad, M Vali, B M Doekemeijer, and J W van Wingerden. A control-oriented dynamic wind farm flow model: WFSim. *Journal of Physics: Conference Series*, 753(3):032005, 2016. URL <http://doi.org/10.1088/1742-6596/753/3/032005>

