

A comprehensive approach to closed-loop wind farm flow control using FLORIDyn

Becker, M.

DOI

[10.4233/uuid:09d71ed3-e5ef-4715-9ca1-ddb73b575221](https://doi.org/10.4233/uuid:09d71ed3-e5ef-4715-9ca1-ddb73b575221)

Publication date

2025

Document Version

Final published version

Citation (APA)

Becker, M. (2025). *A comprehensive approach to closed-loop wind farm flow control using FLORIDyn*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:09d71ed3-e5ef-4715-9ca1-ddb73b575221>

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.



A comprehensive approach
to closed-loop wind farm
flow control using FLORIDyn

Marcus Becker

A comprehensive approach to closed-loop wind farm flow control using FLORIDyn

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof.dr.ir. T.H.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on
Wednesday 2 July 2025 at 17:30

by

Marcus BECKER

Master of Science in Mechatronics,
Technische Universität Darmstadt, Germany
born in Freiburg im Breisgau, Germany

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. J.W. van Wingerden,	
	Delft University of Technology, <i>promotor</i>
Dr. D.J.N. Allaerts [†] ,	Delft University of Technology, <i>copromotor</i>

Independent members:

Prof. dr. S.J. Watson,	Delft University of Technology
Prof. dr. ir. J. Meyers,	Katholieke Universiteit Leuven, Belgium
Prof. dr. M.F. Howland,	Massachusetts Institute of Technology, USA
Prof. dr. ir. L.J. de Vries,	Delft University of Technology
Dr. P.A. Fleming,	National Renewable Energies Laboratory, USA
Dr. ir. S.P. Mulders,	Delft University of Technology, <i>reserve member</i>

Other members:

Dr. ir. B.M. Doekemeijer, Shell, the Netherlands

Dr. ir. B.M. Doekemeijer has, as coauthor of Chapter 2, contributed significantly to the preparation of this dissertation.



Keywords:	closed-loop wind farm flow control, wake steering, yaw-based wake redirection, energy maximization, model-based optimization, nonlinear state-estimation
Printed by:	Gildeprint
Funding:	This work is part of the research programme “Robust closed-loop wake steering for large densely spaced wind farms” with project number 17512, which is (partly) financed by the Dutch Research Council (NWO).
ISBN/EAN:	978-94-6518-093-9

Copyright © 2025 by M. Becker

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

[†]Deceased

All models are wrong, but some are useful.

George E.P. Box

ACKNOWLEDGEMENTS

My PhD journey has been an enjoyable one - driven by curiosity, collaborations, and fun discussions. And while stressful at times, I am privileged to look back on these years as some of the best ones of my life so far. I have to thank a lot of individuals for this privilege and their contributions.

I want to thank Jan-Willem for the guidance and opportunities during my PhD. You trusted me to make my own decisions and gave me possibilities to actively contribute to the research community as an author but also as a task-lead and session chair. Thank you for providing pointers when I was stuck and supporting me when it was needed. I am looking forward to further working together.

Dries, I want to thank you for your curiosity about my research and thorough questions. You helped me enormously to stay critical, strive for better communication, helped me to broaden my horizon, and strengthen my view of the bigger picture. You made me curious about so many topics I did not even know existed. Your input remains valuable and we keep coming back to your advice and work. I will always remember our discussions and time together with a smile.

Thank you to the opposition who agreed to read, judge, and question my work, especially to the international members Michael Howland, Johan Meyers, and Paul Fleming for their travel to Delft or their online availability. Thank you to Simon Watson, Laurens de Vries, and Sebastian Mulders for bringing your independent and outside perspective from within TU Delft. Thank you Bart for being essential to the start and now also to the end of my PhD trajectory. In this context, I also want to thank Daan and Bastian for supervising me during my master's thesis and supporting me in the transition to the PhD. Specifically Bastian for suggesting to think about doing one and establishing the connection to TU Delft, and Daan for picking up the supervision and welcoming me as a new colleague in the office.

Thanks to my two paranimfs Emanuel and Jonas for supporting me during the defense and for the previous collaborations and free time well spent. Emanuel, it was a pleasure to go running, cycling, and climbing with you while chatting about papers and cameras. Jonas, I admire your energy and drive and am very happy that you reached out to write your thesis with me. Thanks for the climbing lessons, the van trip and the exciting reveals of LES results. Generating new ideas with the two of you and exploring them together has been a pleasure. Also thanks to my coauthors - Maarten, as trusty research companion and discussion partner; Maxime, it was a pleasure to keep the Louvain-Delft connection up; Daniel, I am really happy that the seemingly slightly grumpy Dutchman in my office became a long lasting office friend; Leendert for picking up what Daniel, Zhaochen and I left and making it your own; Rafael for being so open and staying in touch from my first conference onwards; Vinit for the fruitful collaboration; Claudia for the successful joint supervision and software ideations; Sebastiano, for supporting new ideas from minute 1 and putting in so much effort.

Further thanks to my wind colleagues, who made coming to the office a bliss. David for the welcoming and curious vibe, it was always fun to exchange new ideas and go climbing together; Mehtab for the great company during fun and difficult times, your open mind, and your willingness to help; Robin, Matteo, and

Tim for holding up the more-than-three-turbines WFFC research torch and having fun discussions; Amr for the warm hugs and amazing wedding experience in Morocco; Unai for the pleasant surprise visits and outdoor activities; Livia, Atin, and Jesse for your keen eyes on figures and for being design partners in crime; Alexandra for accompanying me as TA during NST and now as a colleague; Sachin and Faegheh for picking up the cluster wakes project and making it your own; Guido, Marion, Rogier, and Mees for making sure that the vrिमिबो frequency and bitterballen consumption do not drop below a critical limit; Sebas for your open ear and thoughtful advice; Tristan for your wonderful directness and positive energy; Bert for the wiki-support and frisbee breaks; Fritz for proactively finding your spot in our group; and Jean for the estimator discussions. Outside of TU Delft, I want to thank Paul, Chris, Tony, Jason, Misha, Genevieve, and Ganesh for welcoming me to Boulder and for the numerous discussions; Veronica and Jamie for the shared office time, hikes and talks; Irene, Tuhfe, and Christiann for the collaboration during the IEA Wind Task 44.

My work would have also not been possible without the support and time well spent with all the others working at DCSC. This starts with the secretaries Sandra, Francy, Bo, Erica, Marieke, Heleen, Vanessa, Renate, Katherine, and Maaïke, who have helped me to smoothly navigate any bureaucratic obstacle and put so much personal effort into each social event of the department. Thanks also to Wim and Will for supporting my 3d printing ideas and donating your time and some filament when needed. Thank you to the professors at DCSC, in particular to Sander and Martin for their cooperation during the nonlinear-systems theory course, I learned a lot from you; Bart for caring about my well-being when things got stressful; Manon and Manuel for the shared C-2 floor renovation experience; Dimitris, Tom, Marta, and Riccardo for the many small shared moments in the office. Thank you to Lars, Maria, Maria, Pascal, and Bart for the common TA struggles and victories, and Sander in particular for stepping in when it got tight and offering a helping hand. Further thanks to all of the other (former) PhDs at DCSC, just to name a few: Sam, Gianpietro & Filippo for the belay; Afra for the photography talk; Salim for the friendly competitive vibe; Frida, Clara, and Suad for welcoming me to DCSC; Frederiek for the common PhD-birthday; Paul for your magic and for making Livia laugh; Ivo for your ISS-reaching volleyball skills; Roger, Emilio, Lofti, Coen, Frederik, Léonore, and all the others for the talks and common moments in the office. I also want to thank all the master students I had the pleasure to supervise and support; Stijn, Dongdong, Ihab, Mathieu, Zhaochen, Vibhor, Rozemarijn, Hannah, Julienne - I learned a lot from each one of you and hope that I was able to pass something useful on to you.

Outside of work frisbee kept me more than occupied. The sport was a really important balance for me and I have to thank all the other players for the amazing team experiences in the past years. Without naming everyone, I want to highlight some representatives from each team: First from Ars Ludendi - Duc, Janik, Amir, Hannah, and CK - thanks for the many games played together, parties celebrated, designs made, practices we gave and just overall fun hours spent together. You together with others like Timmon, Datum, Gordon, Jojo, etc. make Darmstadt and Poolimate still feel like homecoming for me. At this point, I also have to mention the APP. Thank you to my flat- and housemates who made my time in the 303 unforgettable; thanks Bene and Lara for the community times, Nathalie for the nightlife, Lea for the fun evenings in the living room, and Mischa for the heated discussions, Paul for the balanced dinners, Maxi & Kat for the Corona time,

Hanna for the breakfasts and Treppenhausparty, just to name a few from the 185b. You made it hard for me to leave Darmstadt but made it easy to talk about the amazing time we had together. With Force Elektro, I have to acknowledge Board 29. Lenny, thanks for setting the board up and being an amazing president, Nick for the throwing and training sessions together and all the sheet and doc shenanigans, Dave for spending more time flying horizontal than running on the field and being an amazing PP&B, Kiet for your humor, kindness, and a sharp eye on our finances, Bernie for your wonderful open Canadian character and your engagement - it was so much fun to work, play and party with all of you. I further want to mention Bram for the many practices we gave together, as well as Dilge, Jasper, Elise, and Jasmin. I have only fond memories of the Mixed teams we trained and played in together. Thanks also to Annemarijn and Chéron for our redesign of the Rising High identity, Ella for supporting my Swiss-Draw madness and Francesca for sharing our common PhD pains. As of writing, Remix is my new playground and I want to thank, among all others, Camilo, Joram, Angie, Niké, Ruben, Emma, and Vince for creating and running the team, as well as putting so much personal dedication into it!

Thanks also to the occasional dinners with Roger, Martijn, Linda, and Martine. You have been a constant for me here and I am so glad we reconnected when I returned to Delft. Talking about dinners, I have to say thank you to the train for the weekends spent together, the dinners and the shared construction work - thanks Jasper, Ellissa, Jan, Veerle, Marnix, Esther, Tim, José, Sigurd, Florens, Inge, Gerald, Kristen, Alexandra, and Rico.

Vielen Dank Simon für die Abende auf fremden Planeten, Rouven für die Wanderlust, Jakob für die Schlachten auf dem digitalen Feld - Danke, dass ihr drei mir regelmäßig einen Grund gegeben habt, um Delft zu verlassen, um ein langes Wochenende zusammen zu verbringen. Ich freue mich schon auf das Nächste! Vielen Dank auch an Katha, Jana, Ann-Kris, Pius, Axel, Max, Kristin, Mara & co. Ich freue mich auf jedes Treffen, die Diskussionen und das Anstoßen. Auf dass die Verbindung bestehen bleibt! Vielen Dank auch an Valentin - ss hat mich gefreut, dass wir wieder in Kontakt gekommen sind, und ich bewundere dein Hausprojekt.

Das bringt mich zu meiner Familie - vielen Dank Anne für die Neugierde in Kunst, Restauration und handwerkliches Arbeiten, Danke Martin für das physikalische Interesse und den Einblick in die Optik. Danke Liesel für die vielen Briefe, den Käsekuchen, die Lektionen an der Nähmaschine und die geteilten Erinnerungen. Vielen Dank auch an Nina, Uli, Anton, Felecie und Nelio. Ralf und Barbara, euch beiden vielen Dank für eure Unterstützung und Begleitung in jungen Jahren und die Flaschen Wein und offene Tür in diesen Tagen.

Bedankt Jan en Anita voor de open deur, de heerlijke etentjes en uitjes, de steun bij de verbouwing, de gedichten en de diepgaande discussies bij een glas wijn. Jullie spelen een grote rol in het feit dat Nederland voor mij inmiddels als een tweede thuis voelt. Veel dank ook aan Jelle en Lianne voor de gezellige Sinterklaas avonden, de gezamenlijke etentjes en de hulp bij de verbouwing.

Vielen Dank an meine Eltern, Barbara und Andreas. Mama, dir zu verdanken habe ich die Reiselust, die Freude am Präsentieren und Lehren sowie meine soziale Ader. Papa, dir habe ich die Neugier an technischen und naturwissenschaftlichen Problemen zu verdanken (aufmachen & gucken, was drin ist), die Freude am Fotografieren und Zeichnen. Danke für eure Liebe, dass ihr mir meinen Werdegang möglich gemacht habt und dass ihr mich zu jeder Zeit unterstützt habt.

Elke, vielen Dank für deine Liebe, Aufmerksamkeit, Unterstützung und Geduld. Es ist schwer vorzustellen, wie ich die letzten Jahre ohne dich bewältigt hätte und ich bin so glücklich, dass wir sie zusammen verbracht haben. Vielen Dank für alles, was du getan hast um mir bei meinem PhD zu helfen - ob es ein offenes Ohr war, die Frage "Wann kommst du denn nach Hause?" oder schlichtweg ein Abendessen. Vielleicht kann ich mich ja ein bisschen revanchieren während deines PhDs.

SUMMARY

Our planet is warming up with potentially disastrous consequences. The main cause of this climate change is the increase of greenhouse gases in the atmosphere, which are mainly emitted by burning fossil fuels to generate energy. Therefore, fossil fuels need to be substituted to reduce emissions from the energy sector. Renewable energies offer an alternative with reduced emissions. Among these, wind and solar energy are growing the fastest. This thesis investigates how the wind energy supply can be increased by improving its operational efficiency.

There are several reasons why a wind turbine may not generate its maximum capacity, one of them being its placement. Turbines are often placed in farms, which allows the collective use of infrastructure and minimizes land usage. The downside is that the turbines influence one another: As a turbine extracts energy from the wind, an area with lowered wind speed develops downstream. This area is called wake, and other turbines affected by it will generate less energy.

The ways to address this problem can be split into pre- and post-construction measures. Pre-construction the wind farm layout can be optimized, and post-construction control strategies are needed to operate the wind farm optimally. These strategies fall under the term *wind farm flow control* and aim to manipulate the flow between the turbines to optimize the farm performance. A turbine's wake can be altered by changing the turbine's resistance to the flow or by misaligning the turbine with the wind direction. The former leads to a faster wake recovery, and the latter results in a redirection of the wake, also called wake-steering.

The current state-of-the-art of wind farm flow control is to utilize wake-steering in an open-loop control configuration. To this end, steady-state engineering models of the wake are used to optimize the farm set points offline. This is done for a selection of atmospheric conditions and the set points are stored in a lookup table (LuT). During operation, the flow conditions are used to look up the precomputed turbine set points. A problem with this approach arises as open-loop control assumes a perfect match between the model and the actual conditions in the field. There are reasons why this might not be the case: (i) There is an inevitable modeling error, which creates a mismatch between the model and the reality; (ii) conditions can arise that are offline not accounted for, e.g., time-varying atmospheric conditions or layout changes due to turbine downtime.

These problems can be addressed by closing the loop. In closed-loop control, measurements are used to continuously correct the model and to adapt it to the current state of the true wind farm. Optimal set points are then found based on the current model state. The control strategy can, therefore, react to new conditions. A challenge is that the optimization needs to happen online and requires a way to incorporate sensor data into the model. Previous work has designed closed-loop approaches using the same computationally cheap steady-state models that were previously used for open-loop control. This was achieved by adapting the parameters of the model based on the mismatch between the observed and predicted measurements, like power generated. A core assumption these models make is that the flow is in a never changing steady state. However, flow conditions do change, and the large spacing between turbines leads to

minutes of delay between the control action the upstream turbine takes and the effect that the downstream turbine experiences. The question arises: What could be achieved using dynamic wake models instead of steady-state ones? These can incorporate wake dynamics, which could lead to better decision-making.

This thesis designs a closed-loop model-predictive wind farm flow control strategy based on a dynamic wake model to maximize the energy generated by a wind farm under time-varying conditions. The thesis is comprised of three building blocks: (i) The development of a dynamic wake model, (ii) the derivation of a sensor fusion strategy to identify the state of the flow field, (iii) the composition of a control strategy that uses the model to optimize the control set points. The building blocks are then connected to form the closed-loop control strategy.

The model building is based on the further development of an existing model, which utilizes a steady-state wake model and reintroduces flow dynamics. In the first step, the underlying wake model is substituted by a three-dimensional one, and the formulation is adapted to heterogeneous flow conditions. In the second step, the model is reformulated as a framework that links to an arbitrary wake model. This is done to profit from advancements in the steady-state model development and to significantly decrease the computational cost of the model. In the third step, the dynamic model is compared to a steady-state one in a set of high-fidelity wind farm simulations under time-varying conditions based on field measurements. The results show that the dynamic model does provide a better match with a simulated wind farm.

In the second part of the thesis, a state estimation methodology is introduced. To this end, an ensemble approach is adopted, where the multiple versions of the model are simulated in parallel. The correlation between the ensembles is then used to correct them based on the predicted and measured wind direction and power measurements of the turbines. A byproduct of the ensemble approach is that each estimated state also has an uncertainty based on how much the ensembles agree on its value.

The third part of this thesis investigates the control and optimization problem. This part focuses on the cost function formulation and the behavior it leads to. In a steady-state frame, the delays do not have to be taken into account, but in a dynamic formulation, they become a challenge. We, therefore, propose a cost-function formulation that synchronizes the control actions with their effect at the downstream turbines. This leads to a series of smaller optimization problems instead of one larger one.

The three building blocks of this thesis are then tested in a case study: The closed-loop controller is employed to maximize the energy of a ten-turbine wind farm under time-varying conditions. Both the farm layout and wind direction time series are based on field conditions. The controller generates an overall energy gain of up to 4 % over the baseline using noise-free wind direction measurements. This is on par with the steady-state approach. However, the closed-loop approach is found to be more robust to disturbed wind direction measurements - Where the performance of the steady-state approach decreases to 1.7 % due to the sensor noise; the closed-loop approach still achieves a 2.5 % gain.

The conclusion of the work presented in this thesis is thereby: Closed-loop wind farm flow control based on a dynamic engineering surrogate model leads to a more accurate and robust state estimation of the wind farm flow field but, given no preview, does not necessarily lead to a higher energy generation than what can be achieved with steady-state models.

SAMENVATTING

Onze planeet warmt op met mogelijk desastreuze gevolgen. De belangrijkste oorzaak van deze klimaatverandering is de toename van broeikasgassen in de atmosfeer, die voornamelijk worden uitgestoten door de verbranding van fossiele brandstoffen om energie op te wekken. Daarom moet het gebruik van fossiele brandstoffen worden geminimaliseerd om de uitstoot van de energiesector te verminderen. Hernieuwbare energiebronnen bieden een alternatief met minder uitstoot. Windenergie en zonne-energie groeien het snelst. In dit proefschrift wordt onderzocht hoe de operationele efficiëntie van windenergie kan worden verbeterd.

Er zijn verschillende redenen waarom een windturbine niet zijn maximale capaciteit genereert. Turbines worden vaak in windparken op zee geplaatst, wat collectief gebruik van infrastructuur mogelijk maakt en landgebruik minimaliseert. Het nadeel hierbij is, dat de turbines elkaar beïnvloeden: Als een turbine energie uit de wind haalt, ontstaat er stroomafwaarts een gebied met een lagere windsnelheid. Dit gebied wordt zog genoemd en andere turbines die hierdoor beïnvloed worden, zullen minder energie opwekken.

De manieren om dit probleem aan te pakken kunnen worden opgesplitst in maatregelen vóór en na de bouw. Voor de bouw kan de lay-out van het windpark worden geoptimaliseerd en na de bouw zijn er regelstrategieën nodig om het windpark optimaal te laten functioneren. Deze strategieën vallen onder de term stromingsregeling voor windparken en hebben als doel de stroming tussen de turbines te manipuleren om de prestaties van het park te optimaliseren. Het zog van een turbine kan worden veranderd door de weerstand van de turbine tegen de stroming te veranderen of door de turbine verkeerd uit te lijnen met de windrichting. Het eerste leidt tot een sneller herstel van het zog en het laatste resulteert in een heroriëntatie van het zog, ook wel zogsturen genoemd.

De huidige stand van zaken op het gebied van stroomregeling van windparken voor energiemaximalisatie is het gebruik van zogsturing in een open regelconfiguratie. Hiervoor worden technische modellen voor de stabiele toestand van het zog gebruikt om de instelpunten van het park offline te optimaliseren. Dit wordt gedaan voor een selectie van atmosferische omstandigheden en de instelpunten worden opgeslagen in een opzoektabel (LuT). Tijdens bedrijf worden de stromingscondities gebruikt om de vooraf berekende turbinesetpoints op te zoeken. Er ontstaat een probleem met deze aanpak omdat open-lusregeling uitgaat van een perfecte overeenkomst tussen het model en de werkelijke omstandigheden in het veld. Er zijn redenen waarom dit niet het geval zou kunnen zijn: (i) Er is een onvermijdelijke modelleringsfout, waardoor een mismatch ontstaat tussen het model en de werkelijkheid; (ii) er kunnen omstandigheden optreden waarmee offline geen rekening is gehouden, bijv. atmosferische omstandigheden die in de tijd variëren of veranderingen in de lay-out als gevolg van stilstand van de turbine.

Deze problemen kunnen worden aangepakt door de lus te sluiten. Bij een gesloten regelkring worden metingen gebruikt om het model voortdurend te corrigeren en aan te passen aan de huidige toestand van het werkelijke windpark.

Vervolgens worden optimale instelpunten gevonden op basis van de huidige toestand van het model. De regelstrategie kan dus reageren op nieuwe omstandigheden. Een uitdaging is dat de optimalisatie online moet gebeuren en een manier vereist om sensorgegevens in het model op te nemen. Eerder werk heeft gesloten-lusbenaderingen ontworpen met dezelfde computationeel goedkope steady-state modellen die eerder werden gebruikt voor open-lusregeling. Dit werd bereikt door de parameters van het model aan te passen op basis van de mismatch tussen de waargenomen en voorspelde metingen, zoals het opgewekte vermogen. Een belangrijke aanname van deze modellen is dat de stroming zich in een nooit veranderende steady state bevindt. De stromingsomstandigheden veranderen echter wel en de grote afstand tussen de turbines leidt tot minutenlange vertraging tussen de regelactie van de stroomopwaartse turbine en het effect dat de stroomafwaartse turbine ondervindt. De vraag rijst: Wat kan er bereikt worden door dynamische zogmodellen te gebruiken in plaats van stationaire modellen? Deze kunnen zogdynamica bevatten, wat tot betere besluitvorming zou kunnen leiden.

Dit proefschrift ontwerpt een gesloten-lus modelvoorspellende stromingsbesturingsstrategie voor een windpark gebaseerd op een dynamisch zogmodel om de energie opgewekt door een windpark onder tijdsvariërende omstandigheden te maximaliseren. Het proefschrift bestaat uit drie bouwstenen: (i) de ontwikkeling van een dynamisch zogmodel, (ii) de afleiding van een sensorfusie strategie om de toestand van het stromingsveld te bepalen, (iii) de samenstelling van een regelstrategie die het model gebruikt om de regelpunten te optimaliseren. De bouwstenen worden vervolgens verbonden tot de gesloten regelstrategie.

De modelbouw is gebaseerd op de verdere ontwikkeling van een bestaand model, dat gebruik maakt van een steady-state zogmodel en de stromingsdynamica opnieuw introduceert. In de eerste stap wordt het onderliggende zogmodel vervangen door een driedimensionaal model en wordt de formulering aangepast aan heterogene stromingsomstandigheden. In de tweede stap wordt het model geherformuleerd als een raamwerk dat gekoppeld is aan een willekeurig zogmodel. Dit wordt gedaan om te profiteren van de vooruitgang in de ontwikkeling van het steady-state model en om de rekenkosten van het model aanzienlijk te verlagen. In de derde stap wordt het dynamische model vergeleken met een steady-state model in een reeks high-fidelity windparksimulaties onder tijdsvariërende omstandigheden op basis van veldmetingen. De resultaten laten zien dat het dynamische model beter overeenkomt met een gesimuleerd windpark.

In het tweede deel van het proefschrift wordt een methode voor het schatten van de toestand geïntroduceerd. Hiervoor wordt een ensemblebenadering gebruikt, waarbij meerdere versies van het model parallel worden gesimuleerd. De correlatie tussen de ensembles wordt vervolgens gebruikt om ze te corrigeren op basis van de voorspelde en gemeten windrichting en vermogensmetingen van de turbines. Een bijproduct van de ensemblebenadering is dat elke geschatte toestand ook een onzekerheid heeft, gebaseerd op hoeveel de ensembles het eens zijn over de waarde ervan.

Het derde deel van dit proefschrift onderzoekt het regel- en optimalisatieprobleem. Dit deel richt zich op de formulering van de kostenfunctie en het gedrag waar deze toe leidt. In een steady-state frame hoeft er geen rekening gehouden te worden met de vertragingen, maar in een dynamische formulering worden ze een uitdaging. Daarom stellen we een kostenfunctieformulering voor die de regelacties synchroniseert met hun effect op de stroomafwaartse turbines. Dit leidt tot

een reeks kleinere optimalisatieproblemen in plaats van één groter probleem.

De drie bouwstenen van dit proefschrift worden vervolgens getest in een case study: De gesloten regelaar wordt gebruikt om de energie van een windpark met tien turbines te maximaliseren onder tijdsvariërende omstandigheden. Zowel de lay-out van het park als de tijdreeksen van de windrichting zijn gebaseerd op veldomstandigheden. De regelaar genereert een totale energiewinst tot 4 % ten opzichte van de basislijn met behulp van ruisvrije windrichtingsmetingen. Dit is vergelijkbaar met de steady-state benadering. De gesloten regelmethode blijkt echter robuuster te zijn bij verstoorde windrichtingsmetingen - waar de prestatie van de stationaire methode afneemt tot 1.7 % door sensorruis, behaalt de gesloten regelmethode nog steeds een winst van 2.5 %.


De conclusie van het in dit proefschrift gepresenteerde werk is daarmee: Closed-loop windpark flow control gebaseerd op een dynamisch engineering surrogaat-model leidt tot een nauwkeurigere en robuustere schatting van de toestand van het stromingsveld van het windpark, maar leidt niet noodzakelijkerwijs tot een hogere energieopwekking dan wat kan worden bereikt met steady-state modellen.

CONTENTS

Acknowledgements	vii
Summary	xi
Samenvatting	xiii
1. Introduction	3
1.1. One-and-a-half degrees Celsius	5
1.2. Wind energy	6
1.3. From one to many turbines	8
1.4. Wind farm flow control	9
1.4.1. The actuated wake	9
1.4.2. Flow sensors	11
1.4.3. Linking measurements to control actions	12
1.4.4. Model-based wind farm flow control	13
1.5. This thesis	17
2. The revised FLORIDyn model: Implementation of heterogeneous flow and the Gaussian wake	33
2.1. Introduction	35
2.2. A new parametric dynamic wind farm model	37
2.2.1. The Gaussian FLORIS model	37
2.2.2. The Zone FLORIDyn model	38
2.2.3. Changes to the FLORIDyn approach	39
2.2.4. Including directional dependency and Observation Point propagation	41
2.2.5. Calculation of C_T and C_P	42
2.2.6. Wind field model	44
2.3. Simulation results	45
2.3.1. Three turbine case	45
2.3.2. Nine turbine case	49
2.3.3. Computational Performance	53
2.4. Conclusions and recommendations	54
2.A. Additional plots and aspects of the simulation results	56
2.A.1. Unfiltered difference between yawed and baseline case	56
2.A.2. Averaged velocity in the nine turbine case	56
3. FLORIDyn - A dynamic and flexible framework for real-time wind farm control	59
3.1. Introduction	61
3.2. Development of a generic FLORIDyn framework	62
3.2.1. Propagation of Observation Points and states	62
3.2.2. Extrapolation from OPs and the creation of temporary wind farms	63

3.2.3. Interfaces and the Immersion and Invariance estimator	66
3.2.4. Implementation	66
3.3. Case study	67
3.3.1. Nine turbine case	67
3.3.2. Performance	69
3.4. Conclusion	69
4. A dynamic open-source model to investigate wake dynamics in response to wind farm flow control strategies	73
4.1. Introduction	75
4.2. Model description	77
4.2.1. Prediction: wake and turbine modeling	78
4.2.2. Correction: linking measurements and states	81
4.2.3. Control: state-based decision making	81
4.3. Case study	82
4.3.1. Simulation setup	82
4.3.2. Predicted controller performance	84
4.4. HKN cases that underwent LES	88
4.4.1. Large-eddy simulation	88
4.4.2. Power generated	89
4.4.3. Power signal correlation	92
4.4.4. Power error statistics	94
4.4.5. Energy generated	96
4.4.6. Computational cost	98
4.5. Conclusion	99
5. Ensemble based flow field estimation using the dynamic wind farm model FLORIDyn	103
5.1. Introduction	105
5.2. Materials and Methods	108
5.2.1. Properties of the FLORIDyn approach	108
5.2.2. Ensemble Kalman Filter formulation	111
5.3. Results and discussion	117
5.3.1. Ensemble Kalman Filter and localization parameters	117
5.3.2. FLORIDyn as validation platform	117
5.3.3. SOWFA as validation platform	120
5.4. Conclusion	127
5.A. State space description	128
5.A.1. Simple propagation	128
5.A.2. Weighted propagation	129
6. Time-shifted cost function design	133
6.1. Introduction	135
6.2. Methodology	136
6.2.1. General setup	136
6.2.2. Time delay estimation	136
6.2.3. Cost function assembly	138
6.2.4. Choosing τ_{ah} and τ_{ph}	140
6.3. Results	141
6.3.1. Simulation	141

6.3.2. Simulation results and discussion	142
6.4. Conclusion	146
7. Closed-loop model-predictive wind farm flow control under time-varying inflow using FLORIDyn	149
7.1. Introduction	151
7.2. Methodology	154
7.2.1. Surrogate model	154
7.2.2. State estimation	155
7.2.3. Controller	155
7.2.4. Reference controllers	161
7.3. Simulation methods	163
7.3.1. “True wind farm” setup	163
7.3.2. Tested controllers and measurements	165
7.4. Simulation results	167
7.4.1. Wake steering	167
7.4.2. Turbine performance	168
7.4.3. Farm level performance	169
7.4.4. Computational performance	172
7.5. Conclusions	172
7.A. Parameter tuning	173
8. Conclusions	177
8.1. Conclusions	179
8.1.1. Dynamic wake modeling	179
8.1.2. State estimation	179
8.1.3. Wind farm flow control	180
8.1.4. Overall conclusions	180
8.2. Recommendations	181
8.2.1. Overall recommendations	181
8.2.2. Dynamic wake modeling	181
8.2.3. State and parameter estimation	182
8.2.4. Control and optimization	182
Epilogue	201
Curriculum Vitae	205
List of Publications	207



1

INTRODUCTION

Nine years after the Paris Agreement was signed, the agreed upon 1.5 °C upper limit for global warming has been violated for the first time. To prevent a further escalation of the climate, we need to lower our greenhouse gas emissions, which are the main drivers of climate change. Part of this transition is the substitution of fossil fuels for alternatives, wind energy being one of them. Wind turbines convert kinetic energy from the flow into electricity. By doing so, they slow down the wind behind them. This affects other turbines, lowering their energy generation. This thesis investigates this issue on a farm scale in the presence of wind direction changes. This thesis presents a control scheme for collaborative wind farm flow control to maximize the energy generated by a farm. The control scheme, in return, has the potential to make fossil fuels more obsolete.

1.1. ONE-AND-A-HALF DEGREES CELSIUS

In 2015 the Paris Agreement was formulated during the 21st Conference of the Parties to the United Nations Framework Convention on Climate Change (COP). A legally binding United Nations (UN) treaty set the goal to keep the global mean long-term temperature below 2 °C compared to the pre-industrial average, preferably below 1.5 °C [1]. To achieve this, signing parties must communicate their plans to reduce their greenhouse gas (GHG) emissions to reach the goals of the Paris Agreement. This goal is set as climate change is driven by human emissions of GHG [2].

In 2024 the Paris Agreement is 9 years old, and the 29th COP is concluded. For this occasion, the World Meteorological Organisation has assembled a report on the state of the climate in 2024. The global mean surface air temperature in January – September 2024 was 1.54 ± 0.13 °C above the pre-industrial average [3]. This is above the 1.5 °C agreed upon 9 years prior. While we cannot conclude from a single year that we have failed to meet this goal, it illustrates the immediateness of climate change.

So, where do GHG emissions come from? From all emissions, 64 % stem from CO₂ emissions due to the use of fossil fuels as energy source [4]. Therefore, two fundamental ways exist to lower GHG emissions from energy generation: (i) reduce the amount of energy used, and (ii) substitute GHG-emitting ways of energy generation with alternatives.

Renewable energies offer one such alternative. They are defined by the United Nations as “[...] *energy derived from natural sources that are replenished at a higher rate than they are consumed. Sunlight and wind, for example, are such sources that are constantly being replenished.*”¹ Throughout their lifetime, they lead to significantly lower GHG emissions than their fossil counterparts: Wind power GHG emissions, for instance, vary between 7.8 and 23 g CO₂-eq./kWh for wind turbines². Coal, on the other hand, varies between 751 and 1095 g CO₂-eq./kWh [5]. Due to their historic use since the industrialization, modern fossil-fuel-based energy sources have a technological development advantage that modern renewable energies need to catch up to. The improvements made to the technologies are mirrored in their Levelized Cost of Energy (LCOE), which puts into perspective how expensive the different energy sources are. Figure 1.1 shows the LCOE development from 2010 to 2023 of on- and offshore wind alongside photovoltaic compared to fossil fuels. A conclusion we can draw from the data is that all three ways of generating electricity have significantly improved in terms of economic viability. However, renewable energies still have a low penetration value from the total final energy consumption perspective. Figure 1.2 shows how the total final energy consumption can be split into four sectors: (i) Industry, (ii) buildings, (iii) transport, and (iv) agriculture. The data shows that renewable energies provide a significant share of the energy used in each sector but do not exceed 16.8 % in any of them. The reason behind this is twofold: the sector does not use electricity as the sole form of energy (e.g., heat), and renewable energies do not have the scale yet to provide sufficient energy to substitute fossil energy sources fully.

¹UN.org: What is renewable energy?, accessed 18. November 2024.

²A carbon dioxide equivalent, abbreviated as CO₂-eq. is a metric measure used to compare the emissions from various greenhouse gases. This is done by converting other gases into their equivalent amount of CO₂ with the same global warming potential. Adapted from Eurostat, statistics explained, accessed 19. November 2024




Figure 1.1.: Levelized cost of electricity generated by wind, solar, and hydropower compared to the range of fossil fuels [6].




Figure 1.2.: Total final energy consumption by end-use sector [7].

The unfolding picture shows that renewable energies are maturing to a point where their energy is more advantageous than fossil fuel energy from a financial and environmental standpoint. Still, they only form a small percentage of the used energy. This begs the question of how the renewable energy supply can be increased to substitute fossil sources further. One option is to increase the installed capacity, e.g., erect new wind turbines or build new solar farms. The other option is to improve the capacity factor, which describes the ratio of actual energy produced by a source compared to its maximum. Wind energy, for instance, had a capacity factor of 41 % to 36 % in 2023 for on- and offshore installations, respectively. This is relatively high compared to the capacity factor of solar with 16.2 %, but lower than the 53 % achieved by hydropower [6]. Since the main global contributors of renewable electric energy in 2023 were hydropower (47 %), on- and offshore wind power (26 %), and solar power (18 %) [7], it makes sense to investigate if their capacity factor can be increased. As also seen in Figure 1.1, wind energy is already a promising candidate to supply a significant share of renewable energies, with offshore wind being more expensive. We therefore investigate how the supply of wind energy can be further increased.

1.2. WIND ENERGY

Wind energy works by converting the kinetic energy of the flow into rotational kinetic energy of the rotor into electrical energy. To this end, horizontal axis wind turbines have blades with an aerodynamic profile; see Figure 1.3. As moving air passes along the blade profile, it creates a pressure and suction side, which accelerates the rotor and causes it to spin. This effect results in

the first conversion from kinetic flow energy to rotational kinetic energy. The




Figure 1.3.: Simplified illustration of the flow passing along a wind turbine blade. The turbine also features a wind vane as well as an anemometer to measure wind direction and speed, respectively.

second conversion happens in the generator by electromagnetic induction: The working principle here is that a moving magnetic field will induce a current in a static wire [8]. In practice, the wind turbine rotor is connected to a shaft, which rotates multiple magnetic poles within the generator's stator. The stator houses wires, which capture the induced currents and make them available to the grid. This conversion creates a force that counteracts the movement of the rotor, slowing it down. Put together, the blades cause the wind turbine rotor to speed up while the generator slows it down, which is in balance at a certain rotor speed. This balance is impacted by the resistance the generator creates and the blade pitch, which regulates how the air passes across the blade. Another component that affects this balance is the orientation of the turbine relative to the flow: a fully aligned turbine will have a more consistent flow across its blades. It will, therefore, operate in a more optimal state.

Since we set out to maximize the wind turbine's energy, it is useful to understand what influences this value. The energy generated by a wind turbine is the power it delivers accumulated over time. The aerodynamic power p of a wind turbine in below-rated conditions can be calculated as follows [9, 10]:

$$p = \frac{1}{2} \rho \pi r^2 u^3 c_p, \quad (1.1)$$

where ρ denotes the air density, πr^2 the rotor area, u the wind speed of flow and c_p the power coefficient. The latter combines the aerodynamic effectiveness of the turbine based on the blade pitch and its orientation. To maximize the energy of a wind turbine, we need to maximize its power. Therefore, the question arises: which components of Equation (1.1) can be influenced to maximize the turbine power? The air density can not be altered; the wind speed, however, depends on the turbine's location and height. The rotor area can be set during construction. Furthermore, the characteristics of c_p are subject to aerodynamic

design and operation and can also be influenced. Based on these relations, wind turbines have grown in size. The quadratic relation between turbine radius and power generation creates a strong return-on-investment effect for each meter for which the turbine blades can be extended. However, increased material demands, structural boundaries, and transportation limitations limit this way of maximizing the turbine power.

The way to further increase the generated power is to distribute its generation across multiple turbines and place them in wind farms. This also comes with advantages related to the shared infrastructure, such as cables or transformers, as well as serial production effects and joint maintenance. However, placing wind turbines together comes with its downsides.

1.3. FROM ONE TO MANY TURBINES

As discussed in Section 1.2, wind turbines work by extracting kinetic energy from the flow. As a result, the flow slows down behind the turbine. This area of lower wind speed is called a wake. Wind turbine wakes can span multiple factors of the diameter of the turbine that caused them. They do recover as they start mixing with the surrounding faster air. Wind turbine wakes of upstream turbines can affect downstream turbines in a wind farm. Going back to Equation (1.1), this leads to a decrease Δu in the wind speed experienced by the turbine. Since the power generated depends on u^3 , a slight reduction of u causes much more significant losses. Figure 1.4 shows Equation 1.1 again, extended by the

The diagram shows the power equation $p = \frac{1}{2} \rho \pi r^2 (u - \Delta u)^3 c_p$ with various parts annotated by colored arrows and text:

- Rotor size** (green arrow) points to r^2 .
- Blade design**, **Blade control**, and **Generator control** (blue arrow) point to c_p .
- Tower height** and **Turbine location** (orange arrow) point to u .
- Wind farm layout** and **Wind farm flow control** (red arrow) point to Δu .

Figure 1.4.: Different approaches to maximize the power generated by a wind turbine, indicated by the part of the equation they affect.

reduction in wind speed due to wakes. We can now annotate the different parts of the equation that we can influence and the related strategies to maximize the power generated by a turbine. The losses induced by wakes can be reduced in two ways: The wind farm layout can be optimized to minimize wake losses. This can be achieved by larger spacing and moving the turbines out of each other's wake based on predominant wind directions. An accompanying approach is to use the turbine's degrees of freedom to change its wake shape. This is known as wind farm flow control (WFFC).

1.4. WIND FARM FLOW CONTROL

Wind farm flow control is an umbrella term for control strategies that alter the flow between the turbines and collectively achieve a farm-wide goal [11]. The underlying premise is that an upstream turbine can change its wake to benefit a downstream turbine. Additionally, turbines share sensor information to gain a better collective insight into their surroundings to determine their best course of action.

In this thesis, we motivate the goal of energy maximization, but wind farm flow control can also be applied to other cost functions. One would be to track a reference power signal with the wind farm [12]. This eases the power fluctuations on the grid and makes the farm output more predictable. A second goal is lifetime preservation - this may entail reducing the turbulent impact on other turbines or de-rating damaged turbines and compensating for their loss with other ones [13]. While these topics are occasionally discussed in this thesis, they are not the main topic and will not further be motivated.

1.4.1. THE ACTUATED WAKE

To better understand the possibilities of performing WFFC, a look at the actuators of a turbine is needed. A wind turbine can alter three degrees of freedom: (i) its orientation, (ii) its blade pitch, and (iii) its generator torque. The latter two change how much and where the rotor extracts energy from the flow. The turbine orientation changes how the flow passes along the rotor plane.




Figure 1.5.: Dynamic induction control approaches to control the wake shape in laminar flow compared to the baseline (a). Collective periodic pitching of the blades results in the Pulse (b), pitching the blades with a phase-offset between them results in the Helix (c). The mechanical engineering faculty building of the TU Delft is given as a reference. The turbine and thus wake size is based on the DTU 10 MW turbine [14].

The effects that can be achieved by varying these degrees of freedom can be summarized into two types: (i) wake redirection and (ii) wake shrinking. The latter is mainly achieved by actuating the blade pitch and changing the generator torque. Applied as a static strategy, the blades pitch out of the wind to reduce the energy they extract. This shortens the wake and minimizes the impact on a second turbine. This control approach is called axial-induction-based control or static induction control [11]. Static induction control has shown to be useful in field tests in a tightly spaced turbine configuration with ≤ 4 turbine diameters distance with full alignment between the turbines [15, 16]. Successive field tests further support the findings but also show that it is difficult to prove consistent gains using static induction control methods [17]. However, as static induction control lowers the force of the turbine against the flow, also the blockage layer changes and decreases [18]. This might lead to a case where data from previous field experiments might underestimate the efficiency of the control method. There is also an opposing approach, where the turbine force onto the flow is deliberately increased [19]. This causes the wind turbine to behave more like a bluff body and induces oscillating motions in the wake, which can increase its recovery [20]. This has made way for a more dynamic approach to induction control: By altering the thrust of the turbines dynamically, the wake mixing with its environment may be enhanced, which leads to a faster recovery and a higher amount of power generated [21]. These results were based on an optimization framework that has full knowledge of the state of the flow, which makes it unfeasible for real application. The resulting actuation signal was therefore reformulated as a sinusoidal variation of the turbine thrust, which can be applied as an open-loop WFFC strategy [22]. The change in thrust can be achieved by collectively pitching the blades. The resulting pulsating wake is illustrated in Figure 1.5 (b) alongside the baseline wake (a) and the faculty building of Mechanical Engineering at the TU Delft³. If the periodic blade pitching is applied with a phase offset between the blades, a helical wake shape emerges [23]. One advantage of this method is that the power generated by the actuated turbine does not fluctuate as much as with the collective change of thrust. The Helix wake is depicted in Figure 1.5 (c). The discussed methods all aim to lower the impact of the wake by either extracting less energy from the flow or by enhancing the wake mixing with the environmental flow. A fundamental assumption often made with these strategies is that there is full alignment between the turbines and that they are relatively densely spaced. If this is not the case, the efficiency of the proposed methods can decrease, and a redirection of the wake becomes a more profitable approach [24].

Redirecting the wake is called wake steering and is achieved by misaligning the rotor with the predominant wind direction [25]: As the air passes the angled rotor plane, a pair of counterrotating vortices forms behind the turbine, leading to a wake deflection, see Figure 1.6 b) and c) [26]. The amount of wake deflection increases with the misalignment angle and the distance to the turbine. The amount of possible wake deflection is large enough to justify wake control for a wide selection of turbine down- and crosswind spacings [24]. This makes it an attractive choice for WFFC applications, and multiple companies offer control solutions based on wake steering [27]. At this point, wake steering is the most advanced approach of WFFC; however, it is still a relatively young approach,

³3DBAG, accessed 26. November 2024.

and during its current design phase, steady-state conditions are assumed. This might yield that control approaches are not fully optimized yet and can be improved. We will, therefore, further investigate steady wake steering as the primary approach to WFFC in this thesis and will explore how it is applied under time-varying atmospheric conditions.




Figure 1.6.: In baseline operation, the wake develops in a straight line (a). If the turbine is misaligned, the wake is deflected to the side (b). This is due to a pair of counter-rotating vortices that curl the wake up and push it to the side (c, adapted from [26]). Figure (a) and (b) show hub-height slices of the wake, (c) a cross-section of (b).

1.4.2. FLOW SENSORS

To determine the appropriate actuation signal, the wind farm controller needs to know the current state of the environment. In its simplest form, this relates to the wind direction and speed at the turbine; in its most complex form, this relates to knowledge about the full turbulent flow field and its future.

The first step is to investigate turbine individual measurements. The power generated, for instance, can be used, along with the rotor speed, to estimate the effective wind speed [28, 29]. Dedicated sensors available with turbines are the wind vane, which indicates the wind direction at the turbine location, and the anemometer to measure the wind speed, see Figure 1.3. The information of the wind speed and direction at the turbine can now be paired with information about the wind farm layout to perform WFFC. However, given the location of the sensors, their measurements are likely affected by the turbine operation and surrounding flow, which leads to signal noise and biases, which are amplified for downstream turbines by the wakes of upstream turbines. One way to combat the noise is to average the value, which simultaneously reduces the amount of data. This measure has led to the adaptation of 10-minute averaged data, which can hide flow dynamics important to wind farm flow control [30]. An alternative is the use of turbine measurements collectively. This can be done by taking the absolute or weighted average [31]. In practice, this does not compensate for possible biases or uncertainties. This has led to the development of more sophisticated approaches, which specifically aim to combat these issues [32, 33]. Another common approach is to use auxiliary measurements from sensors located away from the immediate proximity of the rotor. One tool to achieve this are measurement-masts, as used by, e.g., [16]. These are a collection of flow measurement instruments distributed across different heights and near the turbines but unaffected by their blades and flow. They can provide reference measurements for the turbine but require their own structure to be of a similar

height to the turbines. This is especially disadvantageous and costly in an offshore environment. More advanced sensors exist to measure the flow conditions away from the turbine without a mast. Most prominently, LiDARs (light detection and ranging) have become more specialized and available in wind farms. They can measure the wind speed and direction at defined points in space based on particles flying in the air. Due to their capabilities, they have become a standard asset during measurement campaigns, e.g., [16, 34, 35]. They come with their challenges, such as the amount of cumulated data, the temporal and spatial resolution, the signal-to-noise ratio, and their cost. There has also been work that explores how LiDARs can be used for WFFC purposes [36, 37]. The currently most proven application of LiDAR-assisted control aims to reduce turbine loads [38]. Dual-Doppler radars operate similarly but using a microwave signal instead of a laser. Their temporal and spatial resolution is lower than what can be achieved by a LiDAR; however, they do range further. They have mainly been used as a reference and to analyze wind turbine and farm wakes, e.g., [39, 40], but the required post-processing and uncertainty makes them difficult to use for online decision making [41].

This brief review of available flow sensors shows that the information about the current state of the flow field is limited. Widely available sensors, such as the turbines themselves with their instrumentation, return a coarse and polluted measurement signal of the flow. More advanced sensors like LiDARs and dual-Doppler radars are expensive and difficult to utilize for online decision-making. Additionally they remain the exception and cannot be expected to be available with every wind farm. Based on these findings, a sensor fusion strategy is needed to provide the basis for control decisions. Further, the controller needs to solely rely on the data that is available at runtime.

1.4.3. LINKING MEASUREMENTS TO CONTROL ACTIONS

A wind farm flow control strategy needs to combine both ends: It needs to process and utilize sensor data to find the appropriate reference settings for the available actuators to achieve a farm-wide goal, which, in this thesis, is energy maximization. Linking measurements to control actions can be done in two ways: as a model-free approach or in a model-based manner.

Model-free approaches collect input and output data to identify their relation and utilize it to achieve their goal, e.g. [42–44]. The advantages of this approach are that little to no prior system knowledge is required and that the algorithm can adapt to the exact dynamics at hand. The disadvantage is that the approach takes an increasingly long time to identify the input-output relation, depending on its complexity. Model-based approaches circumvent the learning step by employing a surrogate model of the actual system, in this case, the wind farm and its flow dynamics. This allows the transfer of knowledge about the system dynamics into the control strategy; however, it requires a sufficiently accurate model.

Wind farms possess characteristics that make both model-free and model-based control approaches challenging to apply. First and foremost, the turbine-to-wake-to-turbine interaction is of a highly nonlinear aerodynamic nature. It changes based on the current state of the atmosphere, the surface topography, and inflow conditions. Wind turbines are spaced far apart, which causes delays between a control action and its effect on a downstream turbine.

This can lead to cases where wind direction changes might jeopardize control actions optimized for the initial wind direction. Additionally, the size and financial scale of a wind turbine, and by extension farm, make it difficult to perform experiments on. If experiments are performed, there is no reference wind farm in precisely the same conditions to compare the achieved performance. The scale of the turbines also makes the simulation of wind farms difficult, as simulators need to connect the aerodynamic effects at a blade-profile scale with the flow development across multiple kilometers.

These issues, combined with the limited measurement capabilities discussed in Section 1.4.2, have sparked the development of various wind turbine and wind farm models from high-fidelity numerical codes (e.g., [45, 46]) that allow the closest description of an actual turbine behavior to low-fidelity codes (e.g., [47–49]) that approximate the happenings within a wind farm at a low computational cost. A model-based control strategy can benefit from these developments. In this context, we decide to further investigate possibilities for model-based WFFC, to benefit from the previous model developments, and to adapt them for control purposes.

1.4.4. MODEL-BASED WIND FARM FLOW CONTROL

This form of WFFC needs to make two choices: (i) the surrogate model to represent the wind farm, and (ii) the way control set points are derived and applied at runtime. Both choices are discussed in the following paragraphs.

WAKE MODEL

The choice of the model is essential, as it dictates what kind of control is possible. A more sophisticated model like a Large-Eddy Simulation (LES) coupled with an optimizer can lead to the discovery of new ways of wake actuation [22]. This approach also marks the most demanding approach in terms of computational effort and required knowledge about the flow field. The LES can be down-sampled for more efficient control set point derivation [50], but this approach also requires full knowledge of the flow domain state.

Medium-fidelity wake models, like 2-dimensional Reynolds-averaged Navier–Stokes solvers [51, 52], attempt to capture the flow and wake Dynamics at a reduced computational cost by limiting its dimensionality to a slice at hub height. This can lead to nonphysical effects while applying wake steering, as the wake deflection characteristics are inherently three-dimensional [53]. Free-vortex methods allow for a three-dimensional wake simulation at a low computational cost [54–56]. They solely model the path and interaction of vortex particles shed by the rotor. This method can simulate the wake deflection induced by turbine misalignment, as well as dynamic induction control approaches. However, as the wake breaks down, the approximation tends to become numerically unstable, which makes the method increasingly difficult to use for > 2 turbine simulations and for long-distance wake interaction [57]. Another group of medium fidelity models is described by the Dynamic Wake Meandering (DWM) Model [58]. The approach models the wake as successively shed turbulence boxes that propagate downstream. The transport speed and direction are determined by the present turbulence, which recreates effects such as wake-meandering. Further developments of the method have proven to be

useful for load assessments in a farm context and as medium-fidelity verification of low-fidelity control approaches [59, 60].

Steady-state engineering wake models offer the computationally cheapest approximation of wind turbine wakes, e.g., [47–49]. These derive and fit analytical equations to describe the wake shape, deflection, and reduction of wind speed. Inputs for the models are typically the background wind speed, predominant wind direction, as well as the ambient turbulence intensity. Initially derived for single turbine wakes, wake superposition methods allow the use of the wake models in farms. Their simplicity and speed have made them a common tool to use for tasks like wind farm layout optimization and yield assessments [61–63]. However, the simplifications come at the cost of flexibility and accuracy: The wake models inherently introduce a model error compared to the instantaneous flow based on the steady-state assumption. For control purposes, they can only simulate what they have been designed to represent, e.g., a wake model with no deflection term cannot simulate wake steering. By extension, these models cannot contribute to the discovery of novel WFFC physics, only to the optimization and further development of known ones.

The combination of steady-state wake models with the principles of the DWM model led to the development of dynamic engineering wake models [64]. These induce passive Lagrangian tracers in the flow, which propagate state changes at the rotor plane into the flow field, where they can affect other turbines. Steady-state engineering wake models are then used to calculate the wake impact. This description of a delay dynamic, coupled to engineering wake models, was initially proposed in [64] and has mostly laid dormant since. Given that it describes the mean dynamic wake behavior, it only relies on mean dynamic quantities, such as the predominant wind direction and speed. This substantially lowers the requirements for the sensor needs and makes a derived control strategy more accessible. In this thesis, we, therefore, continue the development of this wake modeling approach for WFFC purposes.

Over the course of this thesis, this type of wake model has had a renaissance. Next to the contributions made in this thesis, several similar models have been published. This namely includes the work of [65], which derives a version that can capture wake meandering by splitting wake and flow field dynamics, [66] which also resorts to modeling flow and wake dynamics apart, [67] who makes the assumption of a tracer inherent flow field state, and [68] which embeds the wake in an LES precursor for the purpose of hardware-in-the-loop experiments. A similar yet fundamentally different approach is presented by [69, 70]: this model presents the wind farm as a network of edges and nodes. The dynamics are achieved by delaying communication between the turbines based on their distance and the atmospheric conditions. The connection between these models and the contributions made in this thesis is discussed in Section 1.5.

CONTROLLER

Model-based WFFC can be applied in an open- or closed-loop configuration, as can be seen in Figure 1.7. The closed- or open-loop distinction relates to how the model is used: in closed-loop, the state of the model is estimated based on measurements. The model and the decision process, therefore, react to the state at hand. In open-loop, the measurements are used to look up control precomputed set points. As a result, there is no remaining flexibility to react to




Figure 1.7.: Combined block diagram of the (a) open- and (b) closed-loop model-based wind farm flow control approaches. The figure is adapted from [11].

states not part of the precomputed database. An advantage of open-loop WFFC is that the control set points can be optimized offline.

Open-loop control has been the predominant form for wind farm flow control, e.g. [35, 71, 72]. The offline optimization is formulated as a steady-state problem to maximize the power generated by the wind farm. Generally, this problem is not convex, which needs to be taken into account during optimization. One tested approach is based on generic algorithms [73], and other approaches use gradient-based information, e.g., [72, 74]. Recently, dedicated algorithms have been developed that derive control set points for wind farms iteratively by testing a broad, coarse range of control set points, which are successively refined [75]. This makes use of the fact that the wake models are relatively cheap to evaluate, while small differences in the final misalignment angle have a small impact on the cost function.

There is little work on the yaw angle optimization using dynamic engineering wake models. One publication uses the dynamic model to perform a grid search to find steady state yaw angles during a synthetic turbine shut-down event in a three-turbine wind farm [76]. A first attempt at closed-loop control has been made using a steady-state wake model [77]. This work continuously identified the wind speed, direction, and wake shape and successively optimized the turbine orientation to maximize the steady-state power generated. Similarly, [78]

also closes the loop on the parameters of the steady-state wake model, which describes its shape. This approach is also presented by [79], in addition to a wind speed estimation.

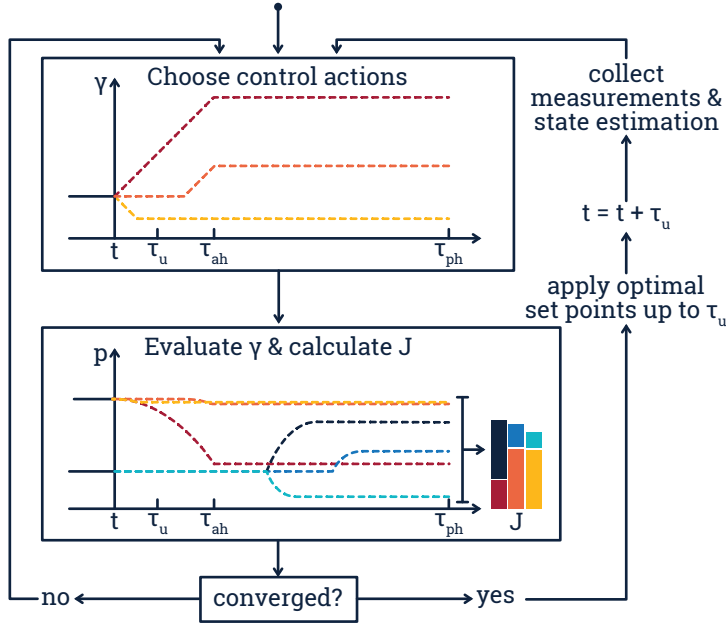


Figure 1.8.: Schematic of economic model predictive control applied to wake steering: first, a set of yaw angle time series is chosen, which vary within the action horizon τ_{ah} . These are then evaluated using the model across the prediction horizon τ_{ph} . A cost function J combines the results into a single number, which is optimized. Once no more improvements are made, the optimal time series is applied, and the time advances by the update time step τ_u . During this time, measurements are taken from the true system, and the state of the model is corrected. Then the optimization is repeated.

A receding horizon approach is used to apply model-based closed-loop control using a dynamic surrogate model. In the case of this thesis, this refers to economic model predictive control (eMPC). In a classic MPC framework, the model would be used to track a reference value. However, in this thesis, we are maximizing the energy and are therefore optimizing towards a supremum. This difference is indicated by the addition of the term economic. This approach is also used by similar work, e.g., [19, 57, 80]. Figure 1.8 depicts the approach adapted to this thesis: based on the current state, a wake model is employed to maximize the energy generated over a prediction horizon using the turbine orientations. The turbine orientation can be varied throughout the action horizon τ_{ah} ; the energy is accumulated up until the end of the prediction horizon τ_{ph} . Once the optimization is done, the control set points are applied to the real system and data is used to estimate the new state of the wake model. After an update time step, the cycle repeats and a new optimization is done.

1.5. THIS THESIS

The state-of-the-art timeline leading up to this thesis shows that there is a gap in wind farm flow control: While model-based steady-state WFFC methods have been tested and refined, dynamic approaches have been limited or non-existent. An initial dynamic engineering wake model has been proposed and preliminary tested, but no further development has been done up to this point. Furthermore, there is little work done on closed-loop wind farm flow control, which may have advantages over its open-loop counterpart. To this end, we formulate an overarching objective:

Thesis objective The development of a closed-loop economic model-predictive control framework using wake steering and a dynamic wake surrogate model to maximize the energy generated by a wind farm under realistic time-varying conditions.




Figure 1.9.: The wake model development contributions and a selection of their immediate scientific context are mostly limited to other dynamic engineering wake models.

To achieve this objective, three building blocks are needed: (i) a dynamic surrogate model, (ii) a state estimation strategy, and (iii) a controller. Thereby the first contribution of this thesis is formulated as:

Contribution I The design and development of a dynamic wake model for wind farm flow control purposes under time-varying conditions.

This entails building upon previous work and further refining the model for control purposes. Figure 1.9 shows the wake model development chapters in their relation to similar wake models. The first contribution is a 3D extension of the previously introduced FLORIDyn model [64]. This model is further adapted for heterogeneous flow conditions and employs a Gaussian wake model. The model is then rewritten as a centerline model, meaning that the wake shape and dynamics are dictated by the wake centerline. This significantly reduces the computational cost. In this iteration, the FLORIDyn model also becomes a framework that can add dynamics to an arbitrary steady-state wake model. In parallel, a number of similar models are published independently [65–67]. Through a collaboration, the last model contribution of this thesis is an open-source object-oriented framework for dynamic parametric wake models like FLORIDyn. It is built in such a way that it can become a central interface to these models.

The second contribution relates to the state estimation of the derived dynamic model:

Contribution II The development of a sensor fusion strategy to identify the state of the flow field based on available turbine data.

This contribution draws from the previous success of the Ensemble Kalman Filter (EnKF) applied to estimate the flow field in a 2D RANS simulator [81], and to do online parameter estimation for a steady-state wake model in a closed-loop configuration [78]. Successive work by [82, 83] also applies this methodology to estimate the wake location.

The last block of the closed-loop controller is the derivation of the turbine inputs. The third contribution is therefore formulated as follows:

Contribution III The derivation of a control strategy that provides control set points in a wind farm to maximize its energy generated.

This contribution needs to adapt a problem that is typically viewed in a steady-state frame into a dynamic one with time-varying atmospheric conditions. To this end, a receding horizon optimization is used, similar to the work of [19, 21, 57, 80]. The contribution further needs to demonstrate how to formulate and solve the optimization problem. Lastly, the three blocks need to be combined and tested for their performance:




Figure 1.10.: State-estimation contribution and its immediate scientific context.

Contribution IV A case study that combines the building blocks and tests the derived closed-loop economic model-predictive wind farm flow control strategy.

The scientific context for the control contributions III and IV is given in Figure 1.11. Note that this is a selection of the relevant literature, which is further extended in the chapters themselves. This work aims to address wind farm flow control in a closed-loop manner under time-varying inflow conditions. This was previously done using steady-state wake models [77, 78, 84], but not yet with dynamic models. During the thesis [57, 80] proposed an adjoint optimization framework for wind farm flow control under time-varying conditions. This work builds on the previous success of adjoint optimization in LES, performed by [19, 21]. The FLORIDyn wake model is not designed for adjoint optimization, which is why a dedicated strategy is developed in contributions III and IV.

THESIS OUTLINE

The derivation of a dynamic wake model is addressed in Chapter 2, 3, and 4: Chapter 2 revises a previously designed dynamic wake model for heterogeneous and changing environmental conditions. The computational load of the model is then greatly reduced in Chapter 3 by a reformulation. This rework of the model also increases the flexibility of the wake model. This is demonstrated in Chapter 4, where the derived model is used to investigate the performance of a




Figure 1.11.: The immediate scientific context of the model-based control contributions in this thesis. Note that this selection either focuses on dynamic models or steady-state models applied in closed-loop.

steady-state controller over a long time horizon. Chapter 5 then addresses the inclusion of measurement data and adapts an ensemble-based state estimation methodology for the previously derived surrogate model. The model-based control approach is then developed and tested in Chapter 6 - 7. Here, Chapter 6 provides an approach that reformulates the energy maximization cost function to neglect the inherent delays between the turbines. Chapter 7 integrates the newly formulated cost function alongside its alternatives in a closed-loop case study. The findings of this thesis are concluded in Chapter 8, which also gives a set of recommendations for future work.

The following chapters 2 to 7 have each been published/submitted as a stand-alone journal or conference paper. Therefore, each chapter will motivate its own work and draw its own conclusions. The related publications are marked with each chapter.




Figure 1.12.: Connection of the thesis chapters to the closed-loop control framework. Chapter 2-4 address the model development, and Chapter 5 proposes a state estimator. Chapter 6 investigates the control problem. Lastly, a case study of the combined closed-loop is presented in Chapter 7.

REFERENCES

- [1] UNFCCC. "Paris Agreement to the United Nations Framework Convention on Climate Change". In: *Paris Climate Change Conference - November 2015*. Paris, Dec. 2015, T.I.A.S. No. 16–1104.
- [2] Intergovernmental Panel On Climate Change (Ipcc). *Climate Change 2021 -- The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 1st ed. Cambridge University Press, July 2023. isbn: 978-1-00-915789-6. doi: 10.1017/9781009157896.
- [3] J. Kennedy, B. Trewin, R. Betts, P. Thorne, P. Forster, P. Siegmund, M. Ziese, S. Mishra, S. Uhlenbrook, J. Alvar-Beltran, A. Gialletti, J. Birner, R.-S. Reusing, E. du Parc, V. Anzellini, S. Ponsérre, V. Grasso, N. Msemu, D. Cuellar Vargas, H. Bastani, R. Boscolo, B. Naran, and C. Ransom. *State of the Climate 2024*. Tech. rep. Geneva: World Meteorological Organization (WMO), 2024.
- [4] Intergovernmental Panel On Climate Change (Ipcc), ed. *Climate Change 2022 - Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 1st ed. Cambridge University Press, Aug. 2023. isbn: 978-1-00-915792-6. doi: 10.1017/9781009157926.
- [5] UNECE. *Carbon Neutrality in the UNECE Region: Integrated Life-cycle Assessment of Electricity Sources*. Tech. rep. Geneva: UNITED NATIONS ECONOMIC COMMISSION FOR EUROPE, 2022.
- [6] IRENA. *Renewable Power Generation Costs in 2023*. Tech. rep. Abu Dhabi: International Renewable Energy Agency, 2024.
- [7] REN21. *Renewables 2024 Global Status Report Collection, Global Overview*. Tech. rep. Paris: REN21 Secretariat, 2024.
- [8] F. T. Ulaby. *Fundamentals of Applied Electromagnetics*. 5. ed. Upper Saddle River, NJ: Pearson/Prentice Hall, 2007. isbn: 978-0-13-241326-8 978-0-13-237139-1.
- [9] W. J. M. Rankine. "On the Mechanical Principles of the Action of Propellers". In: *Transactions of the Institution of Naval Architects* 6 (1865), p. 13.
- [10] R. E. Froude. "On the Part Played in Propulsion by Differences of Fluid Pressure". In: *Transactions of the Institution of Naval Architects* 30 (1889), p. 390.
- [11] J. Meyers, C. Bottasso, K. Dykes, P. Fleming, P. Gebraad, G. Giebel, T. Göçmen, and J.-W. van Wingerden. "Wind Farm Flow Control: Prospects and Challenges". In: *Wind Energy Science* 7.6 (Nov. 2022), pp. 2271–2306. issn: 2366-7451. doi: 10.5194/wes-7-2271-2022.

- [12] E. Ela, V. Gevorgian, P. Fleming, Y. C. Zhang, M. Singh, E. Muljadi, A. Scholbrook, J. Aho, A. Buckspan, L. Pao, V. Singhvi, A. Tuohy, P. Pourbeik, D. Brooks, and N. Bhatt. *Active Power Controls from Wind Power: Bridging the Gaps*. Tech. rep. NREL/TP--5D00-60574, 1117060. Jan. 2014, NREL/TP--5D00-60574, 1117060. doi: [10.2172/1117060](https://doi.org/10.2172/1117060).
- [13] J. Gonzalez Silva, R. Ferrari, and J.-W. Van Wingerden. “Wind Farm Control for Wake-Loss Compensation, Thrust Balancing and Load-Limiting of Turbines”. In: *Renewable Energy* 203 (Feb. 2023), pp. 421–433. issn: 09601481. doi: [10.1016/j.renene.2022.11.113](https://doi.org/10.1016/j.renene.2022.11.113).
- [14] C. Bak, F. Zahle, R. Bitsche, T. Kim, A. Yde, L. C. Henriksen, M. H. Hansen, J. P. A. A. Blasques, M. Gaunaa, and A. Natarajan. “The DTU 10-MW Reference Wind Turbine”. In: May 2013.
- [15] D. Van Der Hoek, S. Kanev, J. Allin, D. Bieniek, and N. Mittelmeier. “Effects of Axial Induction Control on Wind Farm Energy Production - A Field Test”. In: *Renewable Energy* 140 (Sept. 2019), pp. 994–1003. issn: 09601481. doi: [10.1016/j.renene.2019.03.117](https://doi.org/10.1016/j.renene.2019.03.117).
- [16] T. Duc, O. Coupiac, N. Girard, G. Giebel, and T. Göçmen. “Local Turbulence Parameterization Improves the Jensen Wake Model and Its Implementation for Power Optimization of an Operating Wind Farm”. In: *Wind Energy Science* 4.2 (May 2019), pp. 287–302. issn: 2366-7451. doi: [10.5194/wes-4-287-2019](https://doi.org/10.5194/wes-4-287-2019).
- [17] E. Bossanyi and R. Ruisi. “Axial Induction Controller Field Test at Sedini Wind Farm”. In: *Wind Energy Science* 6.2 (Mar. 2021), pp. 389–408. issn: 2366-7451. doi: [10.5194/wes-6-389-2021](https://doi.org/10.5194/wes-6-389-2021).
- [18] E. Bossanyi and J. Bleeg. “How Do Wind Farm Blockage and Axial Induction Control Interact?”. In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092027. issn: 1742-6588, 1742-6596. doi: [10.1088/1742-6596/2767/9/092027](https://doi.org/10.1088/1742-6596/2767/9/092027).
- [19] J. P. Goit and J. Meyers. “Optimal Control of Energy Extraction in Wind-Farm Boundary Layers”. In: *Journal of Fluid Mechanics* 768 (Apr. 2015), pp. 5–50. issn: 0022-1120, 1469-7645. doi: [10.1017/jfm.2015.70](https://doi.org/10.1017/jfm.2015.70).
- [20] L. A. Martínez-Tossas, E. Branlard, K. Shaler, G. Vijayakumar, S. Ananthan, P. Sakievich, and J. Jonkman. “Numerical Investigation of Wind Turbine Wakes under High Thrust Coefficient”. In: *Wind Energy* 25.4 (Apr. 2022), pp. 605–617. issn: 1095-4244, 1099-1824. doi: [10.1002/we.2688](https://doi.org/10.1002/we.2688).
- [21] W. Munters and J. Meyers. “An Optimal Control Framework for Dynamic Induction Control of Wind Farms and Their Interaction with the Atmospheric Boundary Layer”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (Apr. 2017), p. 20160100. issn: 1364-503X, 1471-2962. doi: [10.1098/rsta.2016.0100](https://doi.org/10.1098/rsta.2016.0100).
- [22] W. Munters and J. Meyers. “Towards Practical Dynamic Induction Control of Wind Farms: Analysis of Optimally Controlled Wind-Farm Boundary Layers and Sinusoidal Induction Control of First-Row Turbines”. In: *Wind Energy Science* 3.1 (June 2018), pp. 409–425. issn: 2366-7451. doi: [10.5194/wes-3-409-2018](https://doi.org/10.5194/wes-3-409-2018).

- [23] J. A. Frederik, B. M. Doekemeijer, S. P. Mulders, and J.-W. van Wingerden. "The Helix Approach: Using Dynamic Individual Pitch Control to Enhance Wake Mixing in Wind Farms". In: *Wind Energy* 23.8 (Aug. 2020), pp. 1739–1751. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2513.
- [24] E. Taschner, M. Becker, R. Verzijlbergh, and J. van Wingerden. "Comparison of Helix and Wake Steering Control for Varying Turbine Spacing and Wind Direction". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032023. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032023.
- [25] J. W. Wagenaar, L. Machielse, and J. Schepers. "Controlling Wind in ECN's Scaled Wind Farm". In: *Proc. Europe Premier Wind Energy Event* 01.1 (2012).
- [26] L. A. Martínez-Tossas, J. Annoni, P. A. Fleming, and M. J. Churchfield. "The Aerodynamics of the Curled Wake: A Simplified Model in View of Flow Control". In: *Wind Energy Science* 4.1 (Mar. 2019), pp. 127–138. issn: 2366-7451. doi: 10.5194/wes-4-127-2019.
- [27] S. G. R. Energy. *Siemens Gamesa Now Able to Actively Dictate Wind Flow at Offshore Wind Locations*. Nov. 2019. (Visited on 11/25/2024).
- [28] R. Ortega, F. Mancilla-David, and F. Jaramillo. "A Globally Convergent Wind Speed Estimator for Wind Turbine Systems". In: *International Journal of Adaptive Control and Signal Processing* 27.5 (2013), pp. 413–425. issn: 1099-1115. doi: 10.1002/acs.2319.
- [29] Y. Liu, A. K. Pamososuryo, R. M. G. Ferrari, and J. W. van Wingerden. "The Immersion and Invariance Wind Speed Estimator Revisited and New Results". In: *IEEE Control Systems Letters* 6 (2022), pp. 361–366. issn: 2475-1456. doi: 10.1109/LCSYS.2021.3076040. (Visited on 01/13/2022).
- [30] M. Becker, D. Allaerts, and J. W. van Wingerden. "Wind Pattern Clustering of High Frequent Field Measurements for Dynamic Wind Farm Flow Control". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032028. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032028.
- [31] A. Farrell, J. King, C. Draxl, R. Mudafort, N. Hamilton, C. J. Bay, P. Fleming, and E. Simley. "Design and Analysis of a Wake Model for Spatially Heterogeneous Flow". In: *Wind Energy Science* 6.3 (May 2021), pp. 737–758. issn: 2366-7451. doi: 10.5194/wes-6-737-2021.
- [32] J. Annoni, C. Bay, K. Johnson, E. Dall'Anese, E. Quon, T. Kemper, and P. Fleming. "Wind Direction Estimation Using SCADA Data with Consensus-Based Optimization". In: *Wind Energy Science* 4.2 (June 2019), pp. 355–368. issn: 2366-7451. doi: 10.5194/wes-4-355-2019.
- [33] D. van der Hoek, M. Sinner, E. Simley, L. Pao, and J.-W. van Wingerden. "Estimation of the Ambient Wind Field From Wind Turbine Measurements Using Gaussian Process Regression". In: *2021 American Control Conference (ACC)*. New Orleans, LA, USA: IEEE, May 2021, pp. 558–563. isbn: 978-1-66544-197-1. doi: 10.23919/ACC50511.2021.9483088.


- [34] P. Fleming, J. Annoni, A. Scholbrock, E. Quon, S. Dana, S. Schreck, S. Raach, F. Haizmann, and D. Schlipf. "Full-Scale Field Test of Wake Steering". In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012013. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012013.
- [35] B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann, G. Potenza, F. Calabretta, F. Fioretti, and J.-W. van Wingerden. "Field Experiment for Open-Loop Yaw-Based Wake Steering at a Commercial Onshore Wind Farm in Italy". In: *Wind Energy Science* 6.1 (Jan. 2021), pp. 159–176. issn: 2366-7451. doi: 10.5194/wes-6-159-2021.
- [36] E. A. Bossanyi, A. Kumar, and O. Hugues-Salas. "Wind Turbine Control Applications of Turbine-Mounted LIDAR". In: *Journal of Physics: Conference Series* 555 (Dec. 2014), p. 012011. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/555/1/012011.
- [37] D. Schlipf. *Lidar-Assisted Control Concepts for Wind Turbines*. 1. Auflage. München: Verlag Dr. Hut, 2016. isbn: 978-3-8439-2518-1.
- [38] F. Guo, D. Schlipf, and P. W. Cheng. "Evaluation of Lidar-Assisted Wind Turbine Control under Various Turbulence Characteristics". In: *Wind Energy Science* 8.2 (Feb. 2023), pp. 149–171. issn: 2366-7451. doi: 10.5194/wes-8-149-2023.
- [39] B. D. Hirth, J. L. Schroeder, W. S. Gunter, and J. G. Guynes. "Coupling Doppler Radar-derived Wind Maps with Operational Turbine Data to Document Wind Farm Complex Flows". In: *Wind Energy* 18.3 (Mar. 2015), pp. 529–540. issn: 1095-4244, 1099-1824. doi: 10.1002/we.1701.
- [40] L. Cheung, G. Yalla, K. Brown, N. deVelder, A. Hsieh, T. Herges, D. Houck, D. Maniaci, P. Sakievich, and A. Abraham. "Modification of Wind Turbine Wakes by Large-Scale, Convective Atmospheric Boundary Layer Structures". In: *Journal of Renewable and Sustainable Energy* 16.6 (Nov. 2024), p. 063304. issn: 1941-7012. doi: 10.1063/5.0211722.
- [41] K. Brown, L. Cheung, N. Develder, T. Herges, A. Hsieh, M. Blaylock, G. Yalla, R. Knaus, D. Maniaci, and B. Hirth. "Estimating Uncertainties from Dual-Doppler Radar Measurements of Onshore Wind Plants Using LES". In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092111. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/9/092111.
- [42] J. R. Marden, S. D. Ruben, and L. Y. Pao. "A Model-Free Approach to Wind Farm Control Using Game Theoretic Methods". In: *IEEE Transactions on Control Systems Technology* 21.4 (2013), pp. 1207–1214. doi: 10.1109/TCST.2013.2257780.
- [43] J. Park and K. H. Law. "A Data-Driven, Cooperative Wind Farm Control to Maximize the Total Power Production". In: *Applied Energy* 165 (Mar. 2016), pp. 151–165. issn: 03062619. doi: 10.1016/j.apenergy.2015.11.064.

- [44] P. Hulsman, M. Howland, T. Göçmen, and V. Petrović. “Assessing Closed-Loop Data-Driven Wind Farm Control Strategies within a Wind Tunnel”. In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032049. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032049.
- [45] 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”. In: *Journal of Turbulence* 13 (Jan. 2012), N14. issn: 1468-5248. doi: 10.1080/14685248.2012.668191.
- [46] P. Chatelain, S. Backaert, G. Winckelmans, and S. Kern. “Large Eddy Simulation of Wind Turbine Wakes”. In: *Flow, Turbulence and Combustion* 91.3 (Oct. 2013), pp. 587–605. issn: 1386-6184, 1573-1987. doi: 10.1007/s10494-013-9474-8.
- [47] N. O. Jensen. *A Note on Wind Generator Interaction*. Roskilde, Denmark: Risø National Laboratory, 1983. isbn: 978-87-550-0971-4.
- [48] 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: *2014 American Control Conference*. June 2014, pp. 3128–3134. doi: 10.1109/ACC.2014.6859118.
- [49] M. Bastankhah and F. Porté-Agel. “Experimental and Theoretical Study of Wind Turbine Wakes in Yawed Conditions”. In: *Journal of Fluid Mechanics* 806 (Nov. 2016), pp. 506–541. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2016.595.
- [50] N. Janssens and J. Meyers. “Towards Real-Time Optimal Control of Wind Farms Using Large-Eddy Simulations”. In: *Wind Energy Science* 9.1 (Jan. 2024), pp. 65–95. issn: 2366-7451. doi: 10.5194/wes-9-65-2024.
- [51] A. Rott, S. Boersma, J.-W. van Wingerden, and M. Kühn. “Dynamic Flow Model for Real-Time Application in Wind Farm Control”. In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012039. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012039.
- [52] S. Boersma, B. Doekemeijer, M. Vali, J. Meyers, and J.-W. van Wingerden. “A Control-Oriented Dynamic Wind Farm Model: WFSim”. In: *Wind Energy Science* 3.1 (Mar. 2018), pp. 75–95. issn: 2366-7451. doi: 10.5194/wes-3-75-2018.
- [53] M. J. van den Broek, B. Sanderse, and J.-W. van Wingerden. “Flow Modelling for Wind Farm Control: 2D vs. 3D”. In: *Journal of Physics: Conference Series* 2265.3 (May 2022), p. 032086. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/3/032086.
- [54] Y. Marichal, I. De Visscher, P. Chatelain, and G. Winckelmans. “Towards Physics-Based Operational Modeling of the Unsteady Wind Turbine Response to Atmospheric and Wake-Induced Turbulence”. In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012030. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012030.
- [55] D. Marten. *QBlade: A Modern Tool for the Aeroelastic Simulation of Wind Turbines*. 2020. doi: 10.14279/DEPOSITONCE-10646.

- [56] M. J. van den Broek, D. De Tavernier, P. Hulsman, D. van der Hoek, B. Sanderse, and J.-W. van Wingerden. “Free-Vortex Models for Wind Turbine Wakes under Yaw Misalignment -- a Validation Study on Far-Wake Effects”. In: *Wind Energy Science* 8.12 (Dec. 2023), pp. 1909–1925. issn: 2366-7451. doi: 10.5194/wes-8-1909-2023.
- [57] M. J. van den Broek, M. Becker, B. Sanderse, and J.-W. van Wingerden. “Dynamic Wind Farm Flow Control Using Free-Vortex Wake Models”. In: *Wind Energy Science* 9.3 (Mar. 2024), pp. 721–740. issn: 2366-7451. doi: 10.5194/wes-9-721-2024.
- [58] G. C. Larsen, ed. *Dynamic Wake Meandering Modeling*. Risø R, Report 1607. Roskilde: Risø National Laboratory, 2007. isbn: 978-87-550-3602-4.
- [59] J. M. Jonkman, J. Annoni, G. Hayman, B. Jonkman, and A. Purkayastha. “Development of FAST.Farm: A New Multi-Physics Engineering Tool for Wind-Farm Design and Analysis”. In: *35th Wind Energy Symposium*. Grapevine, Texas: American Institute of Aeronautics and Astronautics, Jan. 2017. isbn: 978-1-62410-456-5. doi: 10.2514/6.2017-0454.
- [60] J. Liew, T. Göçmen, A. W. H. Lio, and G. C. Larsen. “Extending the Dynamic Wake Meandering Model in HAWC2Farm: A Comparison with Field Measurements at the Lillgrund Wind Farm”. In: *Wind Energy Science* 8.9 (Sept. 2023), pp. 1387–1402. issn: 2366-7451. doi: 10.5194/wes-8-1387-2023.
- [61] NREL. “FLORIS. Version 2.1.1”. In: *GitHub repository* (Nov. 2020).
- [62] M. M. Pedersen, A. Meyer Forsting, P. van der Laan, R. Riva, L. A. Alcayaga Román, J. Criado Risco, M. Friis-Møller, J. Quick, J. P. Schøler Christiansen, R. Valotta Rodrigues, B. T. Olsen, and P.-E. Réthoré. “PyWake 2.5.0: An Open-Source Wind Farm Simulation Tool”. In: *DTU Wind, Technical University of Denmark* (Feb. 2023).
- [63] J. Schmidt, L. Vollmer, M. Dörenkämper, and B. Stoevesandt. “FOXES: Farm Optimization and eXtended Yield EvaluationSoftware”. In: *Journal of Open Source Software* 8.86 (June 2023), p. 5464. issn: 2475-9066. doi: 10.21105/joss.05464.
- [64] P. M. O. Gebraad and J. W. van Wingerden. “A Control-Oriented Dynamic Model for Wakes in Wind Plants”. In: *Journal of Physics: Conference Series* 524 (June 2014), p. 012186. issn: 1742-6596. doi: 10.1088/1742-6596/524/1/012186.
- [65] M. Lejeune, M. Moens, and P. Chatelain. “A Meandering-Capturing Wake Model Coupled to Rotor-Based Flow-Sensing for Operational Wind Farm Flow Prediction”. In: *Frontiers in Energy Research* 10 (July 2022), p. 884068. issn: 2296-598X. doi: 10.3389/fenrg.2022.884068.
- [66] R. Braunbehrens, J. Schreiber, and C. L. Bottasso. “Application of an Open-Loop Dynamic Wake Model with High-Frequency SCADA Data”. In: *Journal of Physics: Conference Series* 2265.2 (May 2022), p. 022031. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/2/022031.

- [67] B. Foloppe, W. Munters, S. Buckingham, L. Vandevelde, and J. van Beeck. "Development of a Dynamic Wake Model Accounting for Wake Advection Delays and Mesoscale Wind Transients". In: *Journal of Physics: Conference Series* 2265.2 (May 2022), p. 022055. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/2/022055.
- [68] V. Kipke and C. Sourkounis. "Three-Dimensional Dynamic Wake Model for Real-Time Wind Farm Simulation". In: *2024 32nd Mediterranean Conference on Control and Automation (MED)*. Chania - Crete, Greece: IEEE, June 2024, pp. 808–815. isbn: 9798350395440. doi: 10.1109/MED61351.2024.10566140.
- [69] G. M. Starke, P. Stanfel, C. Meneveau, D. F. Gayme, and J. King. "Network Based Estimation of Wind Farm Power and Velocity Data under Changing Wind Direction". In: *2021 American Control Conference (ACC)*. New Orleans, LA, USA: IEEE, May 2021, pp. 1803–1810. isbn: 978-1-66544-197-1. doi: 10.23919/ACC50511.2021.9483060.
- [70] G. M. Starke, C. Meneveau, J. R. King, and D. F. Gayme. "A Dynamic Model of Wind Turbine Yaw for Active Farm Control". In: *Wind Energy* (Dec. 2023), we.2884. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2884.
- [71] P. Fleming, J. 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". In: *Wind Energy Science* 2.1 (May 2017), pp. 229–239. issn: 2366-7451. doi: 10.5194/wes-2-229-2017.
- [72] M. F. Howland, S. K. Lele, and J. O. Dabiri. "Wind Farm Power Optimization through Wake Steering". In: *Proceedings of the National Academy of Sciences* 116.29 (July 2019), pp. 14495–14500. issn: 0027-8424, 1091-6490. doi: 10.1073/pnas.1903680116.
- [73] 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". In: *Wind Energy* 19.1 (2016), pp. 95–114. doi: 10.1002/we.1822.
- [74] P. Gebraad, J. J. Thomas, A. Ning, P. Fleming, and K. Dykes. "Maximization of the Annual Energy Production of Wind Power Plants by Optimization of Layout and Yaw-Based Wake Control: Maximization of Wind Plant AEP by Optimization of Layout and Wake Control". In: *Wind Energy* 20.1 (Jan. 2017), pp. 97–107. issn: 10954244. doi: 10.1002/we.1993.
- [75] P. A. Fleming, A. P. J. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doekemeijer, and R. Mudafort. "Serial-Refine Method for Fast Wake-Steering Yaw Optimization". In: *Journal of Physics: Conference Series* 2265.3 (May 2022), p. 032109. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/3/032109.
- [76] 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: *2015 American Control Conference (ACC)*. Chicago, Illinois, July 2015, pp. 1702–1708. doi: 10.1109/ACC.2015.7170978.

- [77] B. M. Doekemeijer, D. van der Hoek, and J. van Wingerden. "Closed-Loop Model-Based Wind Farm Control Using FLORIS under Time-Varying Inflow Conditions". In: *Renewable Energy* 156 (Aug. 2020), pp. 719–730. issn: 09601481. doi: 10.1016/j.renene.2020.04.007.
- [78] M. F. Howland, A. S. Ghate, S. K. Lele, and J. O. Dabiri. "Optimal Closed-Loop Wake Steering -- Part 1: Conventionally Neutral Atmospheric Boundary Layer Conditions". In: *Wind Energy Science* 5.4 (Oct. 2020), pp. 1315–1338. issn: 2366-7443. doi: 10.5194/wes-5-1315-2020.
- [79] I. Sood and J. Meyers. "Development and Validation of a Large Eddy Simulation Based Virtual Environment for Optimal Wind Farm Control". PhD thesis. KU Leuven, May 2023.
- [80] M. J. van den Broek, D. De Tavernier, B. Sanderse, and J.-W. van Wingerden. "Adjoint Optimisation for Wind Farm Flow Control with a Free-Vortex Wake Model". In: *Renewable Energy* (Nov. 2022), S0960148122016226. issn: 09601481. doi: 10.1016/j.renene.2022.10.120.
- [81] B. M. Doekemeijer, S. Boersma, L. Y. Pao, and J. W. van Wingerden. "Ensemble Kalman Filtering for Wind Field Estimation in Wind Farms". In: *2017 American Control Conference (ACC)*. May 2017, pp. 19–24. doi: 10.23919/ACC.2017.7962924.
- [82] R. Braunbehrens, S. Tamaro, and C. L. Bottasso. "Towards the Multi-Scale Kalman Filtering of Dynamic Wake Models: Observing Turbulent Fluctuations and Wake Meandering". In: *Journal of Physics: Conference Series* 2505.1 (May 2023), p. 012044. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2505/1/012044.
- [83] J. Di Cave, R. Braunbehrens, J. Krause, A. Guilloiré, and C. L. Bottasso. "Closed-Loop Coupling of a Dynamic Wake Model with a Wind Inflow Estimator". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032034. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032034.
- [84] M. F. Howland, A. S. Ghate, J. B. Quesada, J. J. Pena Martínez, W. Zhong, F. P. Larrañaga, S. K. Lele, and J. O. Dabiri. "Optimal Closed-Loop Wake Steering -- Part 2: Diurnal Cycle Atmospheric Boundary Layer Conditions". In: *Wind Energy Science* 7.1 (Feb. 2022), pp. 345–365. issn: 2366-7451. doi: 10.5194/wes-7-345-2022.



2

THE REVISED FLORIDYN MODEL: IMPLEMENTATION OF HETEROGENEOUS FLOW AND THE GAUSSIAN WAKE

This chapter is based on the following publication:

[1] M. Becker, B. Ritter, B. Doekemeijer, D. van den Hoek, U. Konigorski, D. Allaerts, and J.W. van Wingerden, *The revised FLORIDyn model: implementation of heterogeneous flow and the Gaussian wake*, Wind Energy Science (2022), 10.5194/wes-7-2163-2022.

In this chapter, a new version of the FLOW Redirection and Induction Dynamics (FLORIDyn) model is presented. The new model uses the three-dimensional parametric Gaussian FLORIS model and can provide dynamic wind farm simulations at a low computational cost under heterogeneous and changing wind conditions.

Both FLORIS and FLORIDyn are parametric models which can be used to simulate wind farms, evaluate controller performance and can serve as a control-oriented model. One central element in which they differ is in their representation of flow dynamics: FLORIS neglects these and provides a computationally very cheap approximation of the mean wind farm flow. FLORIDyn defines a framework which utilizes this low computational cost of FLORIS to simulate basic wake dynamics. This is achieved by creating so-called observation points (OPs) at each time step at the rotor plane which inherit the turbine state.

In this work, we develop the initial FLORIDyn framework further considering multiple aspects. The underlying FLORIS wake model is replaced by a Gaussian wake model. The distribution and characteristics of the OPs are adapted to account for the new parametric model but also to take complex flow conditions into account. To achieve this, a mathematical approach is developed to combine the parametric model and the changing, heterogeneous world conditions and link them with each OP. We also present a computationally lightweight wind field model to allow for a simulation environment in which heterogeneous flow conditions are possible.

FLORIDyn is compared to Simulator for Offshore Wind Farm Applications (SOWFA) simulations in three and nine-turbine cases under static and changing environmental conditions. The results show a good agreement with the timing of the impact of upstream state changes on downstream turbines. They also show a good agreement in terms of how wakes are displaced by wind direction changes and when the resulting velocity deficit is experienced by downstream turbines. A good fit of the mean generated power is ensured by the underlying FLORIS model. In the three-turbine case, FLORIDyn simulates 4 s simulation time in 24.49 ms computational time. The resulting new FLORIDyn model proves to be a computationally attractive and capable tool for modelbased dynamic wind farm control.

2.1. INTRODUCTION

In recent years, the topic of wind farm control has gained traction as renewable energies become more and more relevant for the current and future energy mix. Maximizing the power generated by a wind farm is not a trivial task as the turbine-to-turbine interaction is characterized by the complex flow, large delay times and an ever-changing environment. In order to describe the wind field, parametric steady-state approximations have been developed. These describe the mean behavior of the flow with parametrized analytical expressions rather than differential equations. A first approach was presented by [2], which motivated years later the development of more refined steady-state models, such as the Zone FLORIS model [3]. With these low-computational-cost and easy-to-implement wake descriptions, it is possible to develop a model-based control algorithm. These control strategies have managed to improve the power generated in high-fidelity simulations e.g., [3] and in field experiments [4]. The success of parametric steady-state models opens up the question of whether it is possible to overcome one of their great shortcomings: the lack of dynamics. A lowcomputational-cost dynamic wake description can be used to more accurately describe the wake behavior on smaller timescales, during turbine state changes and during environmental changes. This could lead to more sophisticated control approaches and wind farm analysis methods.

There have been efforts to implement parametric models in a dynamic manner, some of which are described here. For a more in-depth discussion of the current state of the art, the interested reader is referred to the review by [5] and more recently, [6]. In the current literature, we have identified two major trails of publications, which will be briefly discussed below.

The first research trail begins with the Aeolus SimWindFarm Toolbox [7], which is publicly available. The toolbox uses the Jensen model [2], coupled with a dynamic description of the centerline and a wind field grid. The centerline would imitate the wake meandering effect based on passive tracers, traveling with the synthetically generated turbulent wind speed. A number of limitations have been imposed for this toolbox: the mean wind speed and direction are constant, the flow field is calculated in 2D, and the turbines operate with fixed yaw angles. The toolbox has enabled the work of [8], who used the Frandsen multiple wake model [9] and added a description of turbine dynamics to estimate fatigue loads. The model is then used to perform induction control based on lookup tables of the thrust and power coefficients with the goal to redistribute loads. This work later inspired the dynamic wind farm simulator, introduced in [10]. The model adds wake steering to the fatigue load estimation and induction control capabilities. To model the effect of yawing the turbine, the deflection formulation of [11] is used. Based on data from the in-house code Bladed, the author formulates the effect of yaw misalignment on the power coefficient by a polynomial expression based on the blade pitch. The wind field is represented by low- and a high-frequency wind speed variations. The low-frequency variations are correlated across the wind farm and cause wake meandering and advection. The highfrequency part is uncorrelated between the turbines and is superimposed with the wake deficits. Lastly, the wake model is switched to the Ainslie model [12].

A second trail of publications can be found starting with [13], where the authors use the previously mentioned Jensen model and extend it to incorporate the impact of time-varying extraction of kinetic energy of turbines due to induction

control. Assuming a constant wind direction and wind speed, the authors derive a linear approximation of the wake advection velocity based on the laws of momentum conservation and mass conservation. The result is a one-dimensional partial differential equation to describe the dynamic wake behavior. The model neglects possible changes of the wake expansion due to a changing thrust coefficient and also does not incorporate yaw angle changes. In [14], the authors extend their model to also take the effects of yawing into account. Most recently, this approach inspired the development of the Floating Offshore Wind Farm Simulator, published in [15]. The authors extend the momentum conservation equations to incorporate time-varying free-stream wind velocity effects. Additionally, they couple the model to a dynamic description of floating platforms, restricted by mooring lines. The authors closely follow [16] to derive a parametric Gaussian velocity shape for their model.

Alongside the two discussed trails of publications, the dynamic wake meandering (DWM) Model was developed. The DWM model, first presented by [17] and later calibrated and refined by [18], proposes an approach much closer to established CFD methods. The model follows a pseudo-Lagrangian approach and creates turbulence boxes around the wake deficit which is created by the turbine. These boxes are then subject to a synthetic turbulent wind field, which allows the modeling of the wake meandering effect. The DWM model puts a focus on load estimation next to the power generated and simulates the turbine by coupling a CFD actuator disc model with an aeroelastic model. Compared to the other mentioned models, the DWM model presents a synergy of CFD methods with engineering approaches.

Another early attempt to derive a dynamic model from a parametric steady-state model was published by [19] who utilized the just published FLORIS model (FLOW Redirection and Induction in Steady-state, [3]) and created the FLORIDyn model (FLOW Redirection and Induction Dynamics). FLORIDyn creates so-called observation points (OPs) at the rotor plane which travel downstream at hub height with the effective wind speed. Their path follows the zone boundaries described by the FLORIS model. The wake deficit and shape depend on the yaw angle and the induction factor. Changes in these variables travel with the OPs and cause a delayed effect at downstream turbines. The authors derive a state-space representation of the model behavior and validate it in a six-turbine simulation against the high-fidelity large eddy simulation environment SOWFA [20]. The state-space representation is then used to implement a Kalman filter for flow field estimation [21]. The model does have shortcomings: due to the two-dimensional flow, shear and veer effects can not be captured, the simulations only work in one wind direction and they do not capture turbulent effects. Furthermore, due to the way the OPs travel, parts of the wake can overlap and can create a faulty wake representation.

In this paper, we aim to overcome these issues and bring the FLORIDyn approach into a form where it can incorporate heterogeneous and changing flow conditions, wind shear, and added turbulence levels. To achieve these changes, we rework the framework to use a Gaussian FLORIS model [16]. This requires a new formulation of the OP behavior. Due to these changes, the wakes can also incorporate locally different and changing flow conditions, such as wind speed, direction and ambient turbulence intensity. To drive the model, a concept of a wind field model is presented as well. The framework is then compared to the simulation environment SOWFA in three- and nine-turbine cases. Furthermore, in order to

allow for collaboration and extension, the code is published in its entirety [22]. The resulting Gaussian FLORIDyn model is a capable, open-source alternative to the few other existing in-house parametric dynamic models, developed for wind farm control purposes.

The remainder of this paper is organized as follows: Section 2.2 discusses the relevant characteristics of the former FLORIDyn framework and how it is adapted. The simulation results are presented in Section 2.3, which also discusses the computational performance. Section 2.4 concludes the paper and gives recommendations for future work.

2.2. A NEW PARAMETRIC DYNAMIC WIND FARM MODEL

In this Section, the new Gaussian FLORIDyn model is introduced. To prevent confusion, we will refer to the models of Gebraad et al. as the Zone FLORIS model [3] and the Zone FLORIDyn model [19]. The Gaussian model by [16] will be referred to as Gaussian FLORIS model.

As the new Gaussian FLORIDyn model is building upon previous work, Sections 2.2.1 and 2.2.2 briefly introduce the terminology and properties of the underlying Gaussian FLORIS model and the Zone FLORIDyn framework. The novel Gaussian FLORIDyn model makes changes to the Zone FLORIDyn framework. These are discussed in Section 2.2.3. Section 2.2.4 describes how heterogeneous environmental conditions are taken into account. To get the power coefficient (C_p) and the thrust coefficient (C_T) values closer to the validation platform SOWFA, a lookup table was generated (Section 2.2.5). Lastly, a basic wind field model is given in Section 2.2.6. It is built to provide the heterogeneous field conditions to evaluate the FLORIDyn model.

In the wake coordinate system, κ_1 , x_1 describes the downwind direction, y_1 the horizontal crosswind direction and z_1 the vertical crosswind direction (Fig. 2.1). In this coordinate frame, the rotor center is always located at $[0, 0, 0]^T$. This coordinate system is not to be confused with the longitudinal (x_0), latitudinal (y_0) and vertical (z_0) world coordinate system κ_0 .

2.2.1. THE GAUSSIAN FLORIS MODEL

The core of the used Gaussian FLORIS model is based on the work of [16]. This work describes a parametric, three-dimensional wake with a Gaussian-shaped wind speed recovery. As it has been applied and described in previous publications (e.g., [23]), only the basic terminology is introduced here as well as the wake shape. In the present work of this paper, the model has been extended with the calculation of added turbulence as proposed by [24]. The power calculation has been extended by the $\cos(\gamma)^{P_p}$ adaptation to the yaw angle [25] and an efficiency term η for tuning [3]. Figure 2.1 depicts an illustration of the wake with its three areas: the potential core, the near-wake area and the far-wake area. For all areas a reduction factor $r = \Delta u / u_{\text{free}}$ can be calculated, where u_{free} is the free wind speed and Δu is the wind speed deficit. The potential core is a region from jets in a coflow [26]. Here, it is used to approximate the immediate region behind the rotor plane. Within the potential core, r is constant. In the near and far field area r reduces to 0, following a Gaussian shape with the extremum at the centerline or border of the potential core. The recovery rate is based on σ_y and σ_z in the respective crosswind directions. The potential core width is described by $w_{y,pc}$

and $w_{z,pc}$, which continuously decrease for the length of the potential core x_c . Lastly, the deflection δ returns the position of the centerline.

The mentioned variables are dependent on turbine states, such as the thrust coefficient C_T and the yaw angle γ , the ambient turbulence intensity I_0 , and a set of 10 parameters. The parameters adjust wake properties such as the recovery rate, the expansion rate, the sensitivity to added turbulence levels and the influence of the yaw angle. The values of the parameters are listed in Table 2.1 in Section 2.3.




Figure 2.1.: Sketched shape of the wake with the different sections, the deflection and areas of equal relative reduction by the Gaussian shape.

2.2.2. THE ZONE FLORIDYN MODEL

An initial FLORIDyn model was published in [19]. The model is based on the previously published Zone FLORIS model, which approximates the wake shape with three zones: near field, far field and mixing zone [3]. Every zone has a formulation of the velocity recovery in downstream direction. To introduce dynamics, observation points (OPs) are created at the rotor plane at each time step.




Figure 2.2.: Creation and propagation of the OPs: In (1) a set of OPs is created, inherits the turbine state and travels downstream, following the FLORIS wake shape, shown in (2). In (3) the turbine state changed and the new OPs inherit a different state (now colored white) and follow the new, dark indicated, wake shape (4).

The OPs serve the purpose to describe the local FLORIS wake characteristics at their location. To do that, they inherit the turbine states at the time of their creation which are necessary to calculate the FLORIS wake. With time, each OP travels downstream, representing a mass of air traveling in the wind. Their travel path is determined by the borders of the FLORIS wake zones. The speed they travel with is equal to the effective wind speed they represent. Figure 2.2 shows the basic concept. Initial OPs are colored black to stress that they inherited the same state. The OPs created after the yaw step are colored white, showing that their inherited state differs.

With this framework, the steady-state wake represents the known FLORIS wake, but other than in FLORIS, changes propagate through the wake instead of instantly affecting turbines downstream. If, for instance, the yaw angle of the turbine changes, the new generation of OPs will inherit the new angle while old OPs still travel according to the previous angle.

In the case of overlapping wakes, an OP travels into the wake of another turbine. The OP locates the closest up- and downstream OPs from the foreign wake and interpolates their reduction factor at its location. In this model, the resulting reduction of the free wind speed is calculated as follows:

$$u_{\text{eff,OP}}(u_{\text{free,OP}}, r_{\text{own}}, r_{\text{f,OP}}) = u_{\text{free,OP}}(1 - r_{\text{own}}) \underbrace{\prod_{i=1}^{n_T} (1 - r_i)}_{r_{\text{f,OP}}}, \quad (2.1)$$

where $u_{\text{free,OP}}$ is the free wind speed at the OP's location. This wake interaction model could also be exchanged for another formulation. The wind speed reduction r_{own} is based on the OP's own wake and r_i is the interpolated reduction of one of the n_T upwind turbines.

To calculate the effective wind speed at the rotor plane, the model calculates an effective velocity reduction factor r_T for every turbine at every time step. The algorithm combines the reduction of each upstream turbine by a root sum square. Within one wake, the reduction factors of the zones are summed, weighted by the overlapping area with the rotor plane.

2.2.3. CHANGES TO THE FLORIDYN APPROACH

Due to the changed underlying FLORIS model, the FLORIDyn approach needs to be adapted. Specifically, the move to a three-dimensional flow field requires a fitting distribution of the OPs, which is discussed in Section 2.2.3. This opens up the possibility to reformulate the calculation of the effective wind speed at the rotor plane, which is presented in Section 2.2.3. The travel speed of the OPs is addressed in Section 2.2.3. In this Section 2.2.3 we use the wake coordinate system \mathcal{K}_1 , indicated by the lower index 1 (eg. $y_{1,\text{OP}}$). The relation between world and wake coordinate system will be explained in Section 2.2.4.

DISTRIBUTION OF THE OBSERVATION POINTS

By changing the underlying FLORIS model, the travel path of the OPs and their distribution has to be rethought. The Gaussian FLORIS model does not have defined borders, and it is three-dimensional. To cover the crosswind wake area regularly for any number of OPs, an algorithm based on the sunflower distribution

was used [27]. The algorithm returns a relative crosswind coordinate $(v_y, v_z) \in [-0.5, 0.5]$ for a given number of OPs. We used 50 OPs per time step. To cover the majority of the Gaussian wake influence, the wake width was chosen to be $\pm 3\sigma_y$ and $\pm 3\sigma_z$ from the centerline and the potential core. The following equation is used to calculate the position of an OP in the wake coordinate system:

$$y_{1,OP}(v_{y,OP}, \sigma_y, w_{y,pc}, \delta) = v_{y,OP}(6\sigma_y + w_{y,pc}) + \delta, \quad (2.2)$$

$$z_{1,OP}(v_{z,OP}, \sigma_z, w_{z,pc}) = v_{z,OP}(6\sigma_z + w_{z,pc}). \quad (2.3)$$

Note that this model only assumes a horizontal deflection. To add a vertical deflection, due to rotor tilt for instance, Equation (2.3) needs to be adapted accordingly. For simplicity's sake σ_y is used, which represents $\sigma_{y,nw}$ for $0 < x_1 \leq x_c$ and $\sigma_{y,fw}$ for $x_1 > x_c$. Respectively, σ_z is defined the same way. The variable δ describes the deflection of the centerline. If OPs travel below $z_1 = 0$ they are ignored. Since v_y and v_z are not changed during the simulation, they can be calculated a priori. They are then used in every time step for the new generation of OPs. OPs with the same relative coordinate follow each other and form what is called a chain. The number of chains is equal to the number of OPs created at each time step.

WIND SPEED AT THE ROTOR PLANE

Since OPs are created at the rotor plane and they interact with foreign wakes, they can be used to estimate the effective wind speed for the power generation. To do that, they have to be distributed across the rotor plane rather than the wake area:

$$y_{1,OP}(v_{y,OP}, \gamma) |_{x_1=0} = v_{y,OP} D \cos \gamma, \quad (2.4)$$

$$z_{1,OP}(v_{y,OP}) |_{x_1=0} = v_{z,OP} D. \quad (2.5)$$

The next step is to determine the area represented by every OP. This is done offline by generating a Voronoi pattern ([28], [29]) with the OPs' relative location as seeds and a circular boundary with radius 0.5. The area of the resulting polygons is normalized by the rotor area and used as weight. All weights are stored in the vector \tilde{w} .

During the simulation, the OPs calculate the reduction of foreign wakes $r_{f,OP}$ on themselves as shown in Equation (2.1). Stored in a vector $\tilde{r}_f = [r_{f,1}, \dots, r_{f,nOP}]^T$ the effective wind speed at the rotor plane is calculated as follows:

$$u_{eff} = \tilde{w}^T (\tilde{r}_f \circ \tilde{u}), \quad (2.6)$$

where \circ stands for the element-wise multiplication and \tilde{u} represents a vector of the free wind speeds at the locations of the OPs. An OP considers itself influenced by a foreign wake if the closest foreign OP is less than $\frac{1}{4}D$ away. This is an arbitrary chosen threshold to reduce the number of OPs for the interaction interpolation. As the outer wake OPs represent the most recovered sections of the wake, this still results in a smooth influence transition.

TRAVEL SPEED

In the former version of the FLORIDyn model, the OPs travel with the effective wind speed they represent. Regions in the center of the wake with lower effective

wind speeds therefore propagate the changes slower than the outer areas. While this seems an intuitive choice, it leads to problems. Initial simulation results showed that, in comparison to the SOWFA simulation, the effects of a state change arrive noticeably slower in FLORIDyn at downstream turbines. Also, due to the difference in OP travel speed, the outer regions adapt their shape earlier in a downstream location, which leads to overlapping areas with the slow regions, which have not adapted yet. This makes the wake representation not injective anymore: multiple OPs occupy and describe the same space at the same time with varying properties.

In this article, the OPs are assumed to propagate with the speed of the free-stream wind rather than the effective wind speed in accordance with Taylor's frozen turbulence hypothesis [30]. The decision is supported by experimental results from [31] and has also been used by other similar codes, e.g. [7]. This also solves the issue of the overlapping wake areas since neighboring OPs travel at the same speed and follow the same state changes. Another implication of this adaptation is that OPs no longer need to calculate the influence of foreign wakes at every time step. This would be used to determine their effective wind speed and thus how far they travel downstream in one time step. The only OPs which need to calculate the foreign influence are the ones at the rotor plane in order to determine the effective wind speed according to Equation (2.6). These model assumptions also significantly decrease the computational load during the simulation. The downside of the change is that the effects of state changes now arrive too fast and abrupt at downstream turbines, which will be seen and discussed with the simulation results in Section 2.3. In future work, the wake propagation speed could be a tuning parameter which is set depending on atmospheric conditions such as the turbulence intensity for instance [32].

2.2.4. INCLUDING DIRECTIONAL DEPENDENCY AND OBSERVATION POINT PROPAGATION

In this section, we address how the OPs, and therefore the wakes, react to a wind direction change. We assume that a wind direction change only affects the wake orientation and that the wake structure and downstream evolution (as defined by the underlying FLORIS model) can be seen independent from the free-stream behavior. It is therefore possible to split the two aspects into two coordinate systems: the world coordinate system \mathcal{K}_0 and the wake coordinate system \mathcal{K}_1 . The free flow conditions are described in \mathcal{K}_0 , whereas the wake properties are described in \mathcal{K}_1 . An OP links these two coordinate systems.

The underlying FLORIS model is described in \mathcal{K}_1 , where the origin $x_1 = y_1 = z_1 = 0$ is located in the center of the rotor plane. The downwind distance is denoted as x_1 , y_1 describes the horizontal crosswind distance and z_1 the vertical one. \mathcal{K}_0 does not have a special orientation apart from $z_0 = 0$ being the ground level and the z_0 axis pointing upwards. In this work, x_0 describes the west-east axis, y_0 the south-north axis. To transform a location vector \vec{r}_1 , described in \mathcal{K}_1 of a turbine with the rotor-center location \vec{r}_0 , into \vec{r}_0 the rotational matrix \mathbf{R}_{01} is used:

$$\vec{r}_0 = \begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} = \vec{r}_0 + \mathbf{R}_{01}(\varphi) \vec{r}_1 = \begin{bmatrix} x_{0,T} \\ y_{0,T} \\ z_{0,T} \end{bmatrix} + \begin{bmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix}. \quad (2.7)$$

This equation assumes a uniform wind direction φ at every location. This will

not be the case for the formulation used for the OP propagation later on in Eq. (2.9). Each OP has two location vectors, $\vec{r}_{0,OP}$ and $\vec{r}_{1,OP}$, one for each coordinate system. The OP's position update and its reduction factor is calculated in \mathcal{K}_1 . \mathcal{K}_0 is used to calculate the wake interaction and to determine the wind speed, the wind direction and the ambient turbulence intensity. At the OP's creation, $\vec{r}_{1,OP}$ is determined by the Equations (2.4) and (2.5) for the crosswind coordinates; the downwind coordinate is set to 0. Its world location, $\vec{r}_{0,OP}$, is then determined by Equation (2.7) with the wind direction $\varphi_{0,T}$ at the turbines location. To iterate the location of an OP from time step k to time step $k + 1$ the downwind step is calculated first in \mathcal{K}_1 :

$$x_{1,OP}(k + 1) = x_{1,OP}(k) + u_{OP} \Delta t, \quad (2.8)$$

where Δt is the time step duration and u_{OP} is the magnitude of the wind vector $\vec{u}_{0,OP}$ at the OPs location $\vec{r}_{0,OP}$. The direction will be applied in Eq. (2.9). For the scope of this work, $\vec{u}_{0,OP}$ can only have non-zero components in x_0 and y_0 direction. With $x_{1,OP}(k + 1)$ the new crosswind locations $y_{1,OP}(k + 1)$ and $z_{1,OP}(k + 1)$ can be calculated with the Equations (2.2) and (2.3), respectively. This completes the transition $\vec{r}_{1,OP}(k) \rightarrow \vec{r}_{1,OP}(k + 1)$. Note that only $x_{1,OP}(k)$ is needed to determine the OP's location in \mathcal{K}_1 . At the cost of calculating $y_{1,OP}(k)$ and $z_{1,OP}(k)$ again at each time step, they do not have to be stored as states. To update $\vec{r}_{0,OP}(k)$ the step which the OP took in \mathcal{K}_1 has to be translated into \mathcal{K}_0 :

$$\vec{r}_{0,OP}(k + 1) = \vec{r}_{0,OP}(k) + \mathbf{R}_{01}(\varphi_{0,OP})[\vec{r}_{1,OP}(k + 1) - \vec{r}_{1,OP}(k)], \quad (2.9)$$

where $\varphi_{0,OP}$ is the wind direction at $\vec{r}_{0,OP}(k)$. Note that $\varphi_{0,OP}$ refers to one OP's individual wind direction; other OPs may have different values. This means that each OP propagates on its own and non-uniform wind directions can be simulated. Figure 2.3 shows the OP step in the wake and world coordinate system. In Figure 2.3-1 and 2.3-2 the wind direction is constant, indicated by the arrow left to the y_0 axis. The OP calculates its step in the wake coordinate system (dotted arrow) and updates its location vectors. These are here simplified to \vec{r}_0 and \vec{r}_1 . In 2.3-3 the wind direction changes and the former FLORIS wake description is invalid and greyed out. With the new wind direction $\mathbf{R}_{01}(\varphi_{0,OP})$ is calculated differently. The OP can calculate its step in the wake coordinate system as before, but its translation $\mathcal{K}_1 \rightarrow \mathcal{K}_0$ changed. Note that neither \vec{r}_0 nor \vec{r}_1 are influenced by the changed wind direction. Their magnitude and orientation remain the same in their respective coordinate systems, however, their orientation towards each other changes.

2.2.5. CALCULATION OF C_T AND C_P

The thrust coefficient C_T is often approximated following the actuator disc theory: $C_T(a) = 4a(1 - a)$, where a is the axial induction factor. To circumvent this approximation, simulations or experiments can be used to create lookup tables. Since most equations of the Gaussian FLORIS model are dependent directly on C_T rather than a , we used lookup tables generated in SOWFA to align FLORIDyn's thrust coefficient with what the turbines in the validation environment experience. For completeness, we also use lookup tables for the power coefficient C_P . The tables are generated for the DTU10 MW reference turbine [33]. It has to be added that these tables are generated from a grid of high-fidelity simulations, where




Figure 2.3.: This figure visualizes the working of Equation (2.9), which is applied for each OP individually. In the transition from (1 \rightarrow 2), the position update of an OP in a time step with a constant wind direction is depicted. The transition from (3 \rightarrow 4) shows the position update when the wind direction changes. In this case, the wake coordinate system is rotated around the OP's location to match the new downstream direction. This causes the apparent origin of the wake in the world coordinate system to change, which is visualized by the grey turbine.

the coefficients were read after the simulation converged to a steady state. The tables can, therefore, only approximate the effect of a changing turbine state and changing wind field conditions onto C_T and C_P . Control approaches for axial-induction-based controllers, such as the one presented by [34], successfully use similar lookup tables, which is why we assume these to be sufficient. Nevertheless, an extension for dynamic circumstances would be a valuable addition for future work but is also connected to a significant computational effort.

In the tables, the coefficients are described dependent on the blade pitch angle β and the tip speed ratio $\lambda(\omega, u_{\text{eff}})$, where ω is the angular velocity of the rotor. However, neither FLORIS nor FLORIDyn can provide λ and β . What they can provide is u_{eff} . Combined with the assumption that each turbine follows a greedy control strategy and maximizes $C_P(\lambda, \beta)$ for the given wind, we can formulate the coefficients dependent only on u_{eff} : first, maximize C_P within the physical limitations of the wind turbine for all wind speeds, then use the $\lambda_{P,\text{max}}$ and $\beta_{P,\text{max}}$ to calculate the respective C_T . The resulting curves can be seen in Figure 2.4.




Figure 2.4.: Greedy control settings of the un-yawed 10 MW DTU reference turbine based on the effective rotor wind speed.

Unfortunately, the resulting $C_T(u_{\text{eff}})$ values can get very high, especially for low wind speeds. This conflicts with some FLORIS equations, which comprise the term $\sqrt{1 - C_T}$ and become complex for C_T values above 1. To avoid these issues, $C_T(u_{\text{eff}})$ is limited to its value at the Betz limit: $C_T|_{a=1/3} = 0.8$ [35]. Another complication is the calculation of the added turbulence levels as it is the only equation which requires the axial induction factor. In this case, the calculation of $C_T(a)$ was inverted to determine $a(C_T)$, based on the actuator disc theory, as follows:

$$a = \frac{1}{2}(1 - \sqrt{1 - C_T}). \quad (2.10)$$

Yaw effects on C_T and a are neglected here. In future work this expression could be substituted, for instance by the polynomial approximation of [36]. It extends $a(C_T)$ to C_T values above 1. However, as C_T is limited in this work, this extension is not necessary. The power coefficient is the remaining aspect which was used unaltered from the lookup tables. For the tested wind speeds below 11 ms^{-1} the power coefficient is constant at $C_P = 0.4929$. The effect of γ is approximated by multiplying C_P with $\cos(\gamma)^{p_p}$. For simplicity's sake we assume p_p to be a constant value. This could be extended by the work presented by [37] which takes the presence of other wakes into account. Similarly, [38] presents an adaptation for locally varying wind profiles.

2.2.6. WIND FIELD MODEL

In order to drive the FLORIDyn model, the wind field needs to be able to simulate heterogeneous, changing environmental conditions. The implemented solution is inspired by the work of [23]. The basic assumption is that measurements of the wind field variables are available at certain locations. This could be due to satellite data, lidar measurements, met masts or other sensors. The value of a measurement for the location of an OP is then interpolated between the measurements available. To reduce the computational effort of an interpolation at every time step, a nearest-neighbor interpolation (NNI) is desirable. To get a sufficient resolution of the measurements to justify an NNI, the sparse measurements \vec{m} have to be mapped to dense measurement grid points \vec{m}_g :

$$\vec{m}_g = \mathbf{M}\vec{m}, \quad (2.11)$$

where the matrix \mathbf{M} describes the mapping and can be calculated offline. The i th row in \mathbf{M} describes the percental composition of $\vec{m}_{g,i}$ from \vec{m} . As a result, the sum of every row in \mathbf{M} is equal to 1. This way, a more complex interpolation can be reduced to a matrix multiplication and a NNI at runtime. In this work, a linear interpolation is used to map the measurements to the grid points, which are spaced in a $20 \times 20 \text{ m}$ grid. OPs outside of the grid defined by \vec{m}_g use the closest grid point. This method is also independent from the quantity measured. In this work, the wind speed, the wind direction and the ambient turbulence intensity were interpolated with the presented method.

However, the presented method is only meant for values changing in the x_0 and y_0 direction. The wind speed is the only field measurement which is also changed in the z_0 direction; wind direction and ambient turbulence intensity are assumed to be constant in the vertical direction. Following [23] the power law is applied:

$$u(z_0) = \left(\frac{z_0}{z_{0,m}} \right)^{\alpha_s} u(z_{0,m}), \quad (2.12)$$

where $z_{0,m}$ is the height of the measurement and α_s is the shear coefficient. The shear coefficient approximates the combined effect of atmospheric stability and surface roughness. A small value describes unstable flow conditions. Examples for characteristic α_s values due to surface roughness are 0.11 over water, 0.16 over grass, 0.20 over shrubs, 0.28 over forests and 0.40 over cities [39]. In this work $z_{0,m}$ is equal to the hub height z_h of the turbine.

2.3. SIMULATION RESULTS

In this section, the Gaussian FLORIDyn model is compared to SOWFA with the focus on turbine interaction. Two wind farm layouts are considered for comparison: three consecutive turbines and a nine-turbine cluster arranged in a 3×3 configuration. The DTU 10 MW reference turbine is used for all simulations. Table 2.1 summarizes the FLORIS and FLORIDyn parameters used in the simulations. The FLORIS parameters k_a and k_b are from [40], $k_{f,a}$ to $k_{f,d}$ are set based on FLORISSE_M [41], and α^* , β^* follow the findings of [16]. The efficiency η was tuned based on turbine T0 in the three turbine baseline case; p_p was tuned based on the three turbine yaw case (Section 2.3.1 and 2.3.1 respectively). For FLORIDyn, n_c relates to the number of OP chains per turbine and n_{op} to the number of OPs per chain. The value of n_{op} was set to cover the entire relevant downstream domain of a turbine; n_c was set to maintain a sufficient density of OPs at the location of other turbines. In FLORIDyn, one time step is 4.0 s long.

Table 2.1.: Parameters used in the simulation with the values they influence
FLORIS

Wake expansion		Added turbulence				
k_a	k_b	$k_{f,a}$	$k_{f,b}$	$k_{f,c}$	$k_{f,d}$	
0.38371	0.003678	0.73	0.8325	0.0325	−0.32	
FLORIS			FLORIDyn		Wind	
Potential core		Power		Chains, OPs	Shear	
α^*	β^*	η	p_p	n_c	n_{OP}	α_s
2.32	0.154	0.8572	2.2	50	200	0.08

Table 2.1 also includes the wind shear coefficient, α_s , which was approximated based on the free flow in SOWFA. The inflow boundary conditions for SOWFA are provided by a precursor simulation which simulates a horizontally homogenous, conventionally neutral atmospheric boundary layer including Coriolis effects. The SOWFA settings differ for the three turbine case and the nine turbine case and will be explained in the respective sections.

2.3.1. THREE TURBINE CASE

The three turbines are placed $5D = 892$ m apart from each other in downwind direction. Turbine T0 is located at (608 m, 500 m), and T1 and T2 are at (1500 m, 500 m) and (2392 m, 500 m) respectively. The mean wind speed at hub height is approximately 8.2 ms^{-1} with an ambient turbulence intensity of roughly 6 %. The mean wind direction is constant along the x axis. The full SOWFA flow field domain spans $3 \times 1 \times 1$ km and was simulated with a time step of $\Delta t = 0.04\text{s}$. The base cells of the flow field are $10 \times 10 \times 10$ m. The




Figure 2.5.: Scaled layout of the three-turbine case with the wind direction indicated by an arrow on the left. The 0° , 10° and 20° yaw orientations from T0 are indicated as turbine symbols with the according orientation. The colored background areas indicate the zones of cell refinement.

refinement areas are centered in the domain and have no offset from the ground. The first refinement is $2.4 \times 0.8 \times 0.5$ km with $5 \times 5 \times 5$ m cells, the second one is $2.2 \times 0.6 \times 0.35$ km with $2.5 \times 2.5 \times 2.5$ m cells. Figure 2.5 shows the to-scale layout including the areas of cell refinement. In SOWFA, the turbines are modeled with the built-in Actuator Line Method (ALM) implementation [42].

To give a better idea of the low-frequency, less-turbulent dynamics, the power generated in SOWFA is also presented filtered by a zero-phase second-order low-pass filter. This non-causal filter is added to aid the visual interpretation of the simulation results. The filter has a damping ratio of $d = 0.7$ and a natural frequency of $\omega = 0.03 \text{ s}^{-1}$. This allows for a more equal comparison as FLORIDyn is sampled at a lower frequency and turbulence is only included as a flow field metric.

A regular second-order low-pass filter with the same d and ω is used for the FLORIDyn data. This causal filter visualizes how low-pass filtering would affect the predicted power generated. This could have advantages due to the changes made to the OP travel speed in Section 2.2.3, which can lead to a very abrupt wake interaction, as will be discussed in Section 2.3.1. However, the filter also naturally adds a phase shift to the data, an effect which might not be desired.

Note that the two filters have different purposes: The non-causal SOWFA filter aims to help to interpret the simulation results, while the causal FLORIDyn filter explores if and when the use of a low-pass filter would be advantageous or if it would decrease the quality of the results.

COMPARISON OF THE WIND FARM START-UP AND STEADY-STATE

In Figure 2.6 the power generated by the turbines in FLORIDyn is compared to the SOWFA simulation. The dynamics in this simulation are the turbulent wind field and the settling of the wake. In the unfiltered data, the interaction in FLORIDyn sets in earlier and more abrupt than in SOWFA. This is due to the OPs traveling at the free wind speed, as explained in Section 2.2.3. The slight curvature of the drop at $t \approx 100 \text{ s}$ can be explained by the wind shear: OPs at a lower altitude travel at a slower free wind speed than OPs at a higher altitude and therefore arrive later at the downstream turbine and therefore affect the turbine later. There are two major aspects to address in order to close the gap between the SOWFA and the FLORIDyn start-up: on the one hand, the way state changes propagate through




Figure 2.6.: Wind farm start-up and steady-state comparison of the power generated in SOWFA and in FLORIDyn. The unfiltered data is plotted in the left figure, and the filtered data is on the right. The SOWFA data is filtered with a zero-phase (noncausal) low-pass filter and FLORIDyn with a causal low-pass filter.

the wake; on the other hand, how a downstream turbine reacts to the new wind field. To give an idea of how a change of the latter aspect would influence the plot, Figure 2.6 also shows low-pass-filtered FLORIDyn data in comparison to zero-phase-filtered SOWFA data. The FLORIDyn data align much more with the SOWFA data but still show discrepancies in terms of dynamic response and steady-state quality of the solution. It should be emphasized that the inFLORIDyn-applied filter does not affect the wake, it only adds an artificial dynamic response to the power calculation. This is important when heterogeneous and changing wind directions are taken into account.

The power generated after the wind farm start-up remains steady in FLORIDyn. This is because there are no turbulent wind speed changes in FLORIDyn. In this state, FLORIDyn is equal to the underlying FLORIS model; therefore, errors in this state need to be solved by adapting the FLORIS model. This could be done by parameter tuning, which has only been done partially in this work (η and p_p ; see introduction Section 2.3).

To incorporate the turbulent wind speed changes, at least to a certain degree, an estimation of the wind speed at the turbine location would be necessary. This could be done by including wind speed sensor data or estimating the wind based on the power generated [21] or based on a torque balance equation [43].

A notable aspect of this simulation is the influence of the added turbulence. Because T0 adds turbulence to the wind field, the wake of turbine T1 recovers faster. Turbine T2 thus experiences higher wind speeds and generates more power than it would without the additional turbulence. This effect can also be observed in the SOWFA data. The old Zone FLORIDyn model is not able to capture this effect due to the underlying Zone FLORIS model. It shows how the FLORIDyn framework is inherently dependent on the capabilities of the employed FLORIS model.

COMPARISON DURING A YAW ANGLE CHANGE

In this simulation, the yaw angle γ of turbine T0 is changing from 0° to 20° in steps of 10° , starting at $t = 200$ s and 800 s. The change rate of γ is set

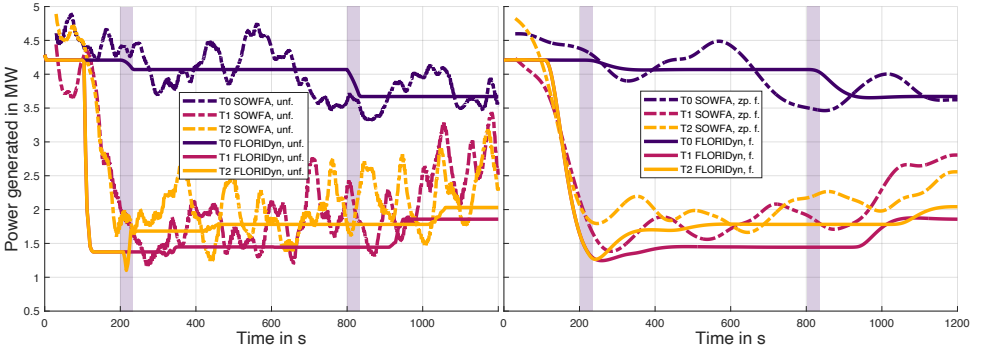


Figure 2.7.: Comparison of the power generated in SOWFA and in FLORIDyn with changing yaw angles. The transparent bars indicate the time window in which turbine T0 increases its yaw angle by 10° . The left plot shows the unfiltered data, and the right one shows the filtered data. The SOWFA data is filtered with a zero-phase (noncausal) low-pass filter, FLORIDyn with a causal low-pass filter.

to 0.3° s^{-1} . Figure 2.7 shows the unfiltered SOWFA data in comparison to the unfiltered FLORIDyn data on the left, as well as the filtered data on the right. Filtering was performed as described in the introduction of Section 2.3.1.

In FLORIDyn, turbine T1 shows a slight reaction to the yaw changes of turbine T0 at roughly $t \approx 320 \text{ s}$ and more significantly at $t \approx 920 \text{ s}$. The influence of the state change then travels further and impacts T2 at $t \approx 430 \text{ s}$ and, as well more significantly, at $t \approx 1030 \text{ s}$. In SOWFA, the reaction is obscured by turbulent influences. However, an increase in average power can be seen for T1 and T2 throughout the entire simulation. Figure 2.8 shows the baseline simulation in comparison to the SOWFA simulation, in absolute values and the difference. The data of both simulations can be compared since they use the same wind field. It allows for a more accurate determination of the reaction time to the upstream change. Turbine T1 starts to react to the first yaw angle change at $t \approx 320 \pm 8 \text{ s}$, T2 at $t \approx 434 \pm 8 \text{ s}$. Given the distance to T0, this translates to a travel speed of the first influence between $[6.98, 7.97] \text{ ms}^{-1}$ to T1 and $[7.38, 7.90] \text{ ms}^{-1}$ to T2. This indicates that first effects of the yaw angle change do travel at almost free-stream velocity and the times align with the FLORIDyn prediction. However, FLORIDyn does lack the dynamic nature of the interaction, which means the response of the wake to the state change of the upstream turbine and the response of the downstream turbine to changes in the wake. Given that all OPs travel at their free-stream velocity, turbine state changes are directly picked up by the OPs and transported, and the FLORIDyn turbine reacts immediately when the OP arrives. The low-pass-filtered results provide an idea of how a dynamic response could change the results. The unfiltered difference between the SOWFA simulations is given in the Appendix 2.A.1. Figure 2.7 also shows that FLORIDyn underestimates the gain in generated energy in the steady-state region. The error likely lies with the underlying FLORIS model as it is a steady-state error. Better parameter tuning would likely decrease or even eliminate the error.




Figure 2.8.: Comparison between the zero phase filtered (zp.f.) SOWFA data in the baseline case (bl.) and in the yaw case. On the left are absolute values, and on the right is the difference between the yaw case and the baseline case. The transparent bars indicate the time window in which turbine T0 increases its yaw angle, first from 0° to 10° , then to 20° .




Figure 2.9.: Complete nine turbine FLORIDyn flow field in comparison to SOWFA at $t = 600$ s. The wind direction change is indicated in the lower left corner of the SOWFA plot.

2.3.2. NINE TURBINE CASE

In order to test the model in a changing environment, a simulation with nine turbines was performed. The turbines are arranged in a three-by-three grid with 900 m distance to the closest turbines and 600 m to the edge. The setup is presented in Figure 2.9 as well as the numbering of the turbines. The wind field performs a 60° uniform wind direction change from 15° to 75° , as indicated in Figure 2.9. The change starts at $t = 600$ s with 0.2° s^{-1} and ends at $t = 900$ s. The change in wind direction is achieved by using SOWFA's built-in utility to specify the wind speed and wind direction at a certain height and time. For the remainder of the simulation, the wind field conditions remain steady. To keep the computational load of the SOWFA simulation low, the DTU 10 MW turbines were simulated with the Actuator Disc Method (ADM). This also allows a coarser grid and time resolution: The domain is discretized in $10 \times 10 \times 10$ m cells, and the SOWFA time step length is set to 0.5 s. The flow field spans $3 \times 3 \times 1$ km. The average wind speed during the simulation is 8.2 ms^{-1} and the ambient turbulence intensity is




Figure 2.10.: Nine turbine flow field at hub height during the wind direction change at $t = 700$ s (top), $t = 800$ s (center) and $t = 900$ s (bottom). The FLORIDyn flow field is on the left and includes grey arrows as an indicator of the current wind direction. On the right is the corresponding snapshot from the SOWFA simulation.

approximately 6 %. During the simulation, the turbines maintain a yaw angle of 0° and turn with the wind. For simplicity we assume ideal wind direction tracking capabilities and apply a prescribed motion. For more information see the dataset which contains the SOWFA files for the case and the precursor simulation [44].

Figure 2.10 shows the flow field during the wind direction change, starting at the time instances $t = 700$ s, $t = 800$ s and $t = 900$ s. The SOWFA slices are taken at hub height. To show the center of the FLORIDyn flow field, only OPs between $0.5 z_h < z_{OP} < 1.5 z_h$ are plotted. As the wake expands, more chains leave these bounds, which leads to a sparser description of the wake in the plot. Due to the changes to FLORIDyn described in Section 2.2.3, the OPs do not influence each




Figure 2.11.: Power generated in the nine-turbine case. The grey area marks the time window in which the wind direction linearly changes by 60° . The plots are arranged to fit the layout of the wind farm in Figure 2.9. The data shows the zero-phase filtered (zp. f.) and the unfiltered (unf.) SOWFA data, as well as the filtered (f.) and unfiltered FLORIDyn data.

other in the field and OPs with a higher velocity can appear among OPs with a lower velocity. The net effect of multiple OPs is only calculated at the rotor plane. The grey arrows in the FLORIDyn plots indicate the current wind direction. The plots of both simulations visualize how the wakes slowly transition to the new wind direction, forming a bow shape in the process. FLORIDyn seems to describe the general path of the SOWFA wakes quite well. It also capture some effects like shorter, wider wakes of T4 and T5 at $t = 900$ s.

To more accurately judge the timing of FLORIDyn, Figure 2.11 shows the generated power of all nine turbines. The plots are arranged in the same way as the turbines in the flow field plots. All plots show the filtered and unfiltered data of FLORIDyn and SOWFA. The filtering is identical to the filtering in Section 2.3.1. The grey area marks the time window of the wind direction change. Looking at the magnitude of the generated power, FLORIDyn predicts the average of the free stream turbines T0, T3, T6, T7 and T8 quite well, but the remaining turbines show

a noticeable offset in generated power. This could be due to speed-up effects and is briefly discussed in Appendix 2.A.2. The interesting aspect is the timing of the wake interaction from upstream turbines with downstream turbines. The generated power by T4 shows the passing of the wake of T6 during the wind direction change. Noticeable is the accuracy with which the unfiltered FLORIDyn data align the unfiltered SOWFA data. Table 2.2 lists the points in time at which the power generated is minimal in SOWFA and in FLORIDyn, as well as the difference. This shows that FLORIDyn predicts the maximal wake influence 5.5 s later than in SOWFA. The filtering of the FLORIDyn data significantly worsens the quality of the

Table 2.2.: Points in time at which the power generated in the nine-turbine case is minimal due to wake interaction.

	Turbine 1		Turbine 2		Turbine 4	Turbine 5
	Min. 1	Min. 2	Min. 1	Min. 2	Min. 1	Min. 1
SOWFA (s)	833.5	996.5	822	972	826.5	809
FLORIDyn (s)	832	992	836	992	832	836
Error (s)	−1.5	−4.5	+14	+20	+5.5	+27

result, in contrast to the simulations of the three-turbine case. The filtering was applied on the data from the rotor plane. Thus, only modifying the way a turbine perceives the incoming, foreign wake in FLORIDyn will not improve the simulation for changing environments. As a result, future research has to improve the way a turbine dynamically influences its own wake. Turbine T1 shows similar behavior to T4: the generated power shows the wake influence of T3 first, but it also shows, after the wind direction stopped changing, the influence of the outskirts of the wake of T6. While the timing of this interaction shows good agreement, the magnitude of the interaction is considerably lower in FLORIDyn than in SOWFA. This could be due to a too fast recovering FLORIS wake, an inadequate wake superposition method or due to local turbulence levels, which FLORIDyn can not capture. T5 shows the overlapping influence of the wakes of T6 and T7. The two overlapping Gaussian influences do form a longer period of reduced generated power. This can be seen in both simulations. Table 2.2 shows the largest timing error between SOWFA and FLORIDyn for this wake influence. This could stem from an inaccurate wake interaction model and the way added turbulence is treated. T2 shows the most overlapping influences by the wakes of T3, T4, T6 and T7 in this order. While the first two overlapping interactions show good agreement, FLORIDyn shows poorer agreement with SOWFA for the influence of T6 and T7. Again, the SOWFA simulation suggests a larger decrease in generated power. A reason could be that the way wakes combine is not described accurately enough. The timing of the wake of T7 seems to be a bit too late as well. However, the SOWFA simulation recovers to its steady-state at about the same time as FLORIDyn. Additionally, all turbines, with the exception of T6, experience a small influence of the upstream turbines in the steady-state configuration. This effect is not noticeable in the SOWFA simulations. Concluding, FLORIDyn describes the timing of passing wake influences quite well. However, there are discrepancies in terms of magnitude and possibly in the way wakes combine their effects on downstream turbines.

2.3.3. COMPUTATIONAL PERFORMANCE

Table 2.3 contains the average computational time per time step, which is equivalent to 4 s simulation time. This can be compared to SOWFA, which can take around $5.8 \cdot 10^2$ s to $5.4 \cdot 10^3$ s per core, per time step, depending on the setup [45]. The FLORIDyn measurements were performed for two and three consecutive turbines and a 2×2 and 3×3 turbines wind farm. The times exclude plotting and the simulation setup time. A setup can take up to 3 s, depending on how much data needs to be imported. The measurements were taken on a MacBook Pro (2019), 2.3 GHz 8-Core Intel i9 CPU, 32GB of 2667 MHz DDR 4 RAM, an SSD and MacOS Catalina (10.15.7). The simulation environment is Matlab 2020a without the use of toolboxes, such as the parallel computing toolbox, and without precompiled code, besides what is built into the simulation environment. These results naturally vary with the layout, atmospheric behavior, simulation settings, etc., and are only meant to give an estimation of the performance.

A first takeaway is that FLORIDyn simulates all cases faster than real-time: Within 4 s, the simulation can perform between 164 and 6.5 simulation steps, depending on the number of turbines simulated. This results in 656 s_{sim} in-simulation time for two turbines and goes down to 26 s_{sim} for nine turbines in 4 s of real-time. This opens up the needed computational headroom for a model-based real-time control strategy and the necessary optimization. On the other hand, the times also do not offer a large time window for optimization. For instance, in the three-turbine yaw case, it takes roughly 300 s_{sim} in simulated time until the yaw changes have propagated from the first turbine to the last turbine. To optimize the control actions with this model for the near future, parallel computing would be needed. With an increasing number of turbines, this time window decreases.

Generally, the computational time increases quadratically with $n_T^2 - n_T$, as n_T turbines need to determine if they are in the wake of another turbine and calculate the influence. This growth in computational effort is assumed to decrease with larger wind farms: As the spacial dimension grows, not every turbine needs to consider all other turbines for interaction. Nevertheless, the simulation times will exceed what is practical for a wind farm with $\gg 3$ turbines.

There are multiple opportunities to improve performance that have not been utilized so far. The main aspect that increases computational effort is the interaction among the turbines. The first step can be to calculate the turbine interactions in parallel. A second step is to find a way to efficiently determine if a turbine is influenced by a wake. Furthermore, the number of OPs per turbine can be decreased and tuned: not all chains need to be equally long, and OPs that wander out of the domain can be disregarded. Then, there is the fundamental question of whether the proposed structure of FLORIDyn can be improved, for instance, by using less but more efficient OPs. Eventually, the programming platform can be switched to a choice that allows more specific optimization, e.g., C, C++, or Julia.

Table 2.3.: Computational performance

Number of turbines	2	3	4	9
Total number of OPs	$2 \cdot 10^4$	$3 \cdot 10^4$	$4 \cdot 10^4$	$9 \cdot 10^4$
$t_{comp. \text{ time per step (s)}}$	$2.44 \cdot 10^{-2}$	$5.87 \cdot 10^{-2}$	$1.09 \cdot 10^{-1}$	$6.13 \cdot 10^{-1}$

2.4. CONCLUSIONS AND RECOMMENDATIONS

In this paper, a new FLORIDyn model is presented and compared to SOWFA simulations. This model utilizes a Gaussian FLORIS model and concepts from the previously published FLORIDyn framework by [19] to create a three-dimensional, dynamic and computationally lightweight wind farm model. The new FLORIDyn model is further capable of simulating its wakes under heterogeneous and time-varying flow conditions. To achieve this, we presented a mathematical approach to decouple wake and flow characteristics into two coordinate system which are connected by observation points. To simulate changing environmental conditions, a method to map sparse flow field measurements to a finer grid was presented which avoids the interpolation cost at runtime. The new FLORIDyn model shows good performance compared to SOWFA in terms of timing and is able to predict accurately when a downstream turbine is experiencing influences from upstream turbines.

Despite the considerable advancements over the old FLORIDyn implementation, there are still several aspect of the model which can be improved. The central aspect is how turbines influence wakes and how wakes are perceived by turbines. In this work we have decoupled the OP propagation speed from the effective wind speed, which effectively leads to a simpler, lightweight model while the wake behavior is still dynamically described. However, this way state changes reach downstream turbines too soon and in a sudden manner. Ideally this can be overcome by finding better, computationally lightweight methods to model the influence of changing turbine states on the wake needs and also how a turbine reacts to dynamic changes in the flow. Another aspect that can be improved is related to the interface between FLORIDyn and FLORIS. FLORIS has been subject to many developments and improvements, and FLORIDyn can utilize these improvements if it improves the interface: with a generic interface, newer developments can be included and existing code can be used in a sustainable manner. The simulations also show that parameter tuning has to be more accessible and possibly needs to be performed online in some cases. The next aspect which could be improved is the coupling of FLORIDyn with the turbulent environment of the real wind farm (or its surrogate). Combined with the changes to the OP propagation speed from this work, this can lead to a more uneven OP distribution with dense areas where high wind speeds decrease and sparse areas where low wind speeds increase. An extension to the model could feature a method to combine and generate OPs, depending on the density of OPs. Although, this could also lead to undesired information loss, depending on the implementation. To achieve better results, the wind field model has to be replaced or enhanced by estimators. The latter would provide a more accurate estimate of wind speed, direction and ambient turbulence intensity for the FLORIDyn simulation. In the long term, a dynamic description of the environment could become part of the FLORIDyn model. This could also include effects like induction zones and speed-up between the turbines. The last aspect to consider for improvement is the performance. In its current implementation, FLORIDyn delivers its results at a low computational cost. This has to be maintained, if not improved, to allow its use for dynamic real-time closed-loop control algorithms in the future. The simulation also needs to be structurally improved to keep its low computational cost for wind farms with large numbers of turbines.

Concluding, the new FLORIDyn model is a promising concept with unique

strengths. With FLORIS in its core, it utilizes an existing, successfully employed model and provides a new dimension in a challenging environment at a low computational cost. The model can already be adapted to work in a closed-loop control design and shows more potential if the mentioned aspects are improved.

2.A. ADDITIONAL PLOTS AND ASPECTS OF THE SIMULATION RESULTS

2.A.1. UNFILTERED DIFFERENCE BETWEEN YAWED AND BASELINE CASE




Figure 2.12.: Difference between the yaw case in SOWFA and the baseline case with unfiltered and zero-phase-filtered data. Filtering was performed before calculating the difference. The dotted lines mark the start of the yaw angle changes of T0.

Figure 2.12 shows the difference between the power generated in SOWFA in the yaw case (Section 2.3.1) and the steady-state base line case (Section 2.3.1). Both simulations are performed in the same turbulent environment, something which would be impossible to achieve in realistic conditions. This way, the difference allows for a clearer interpretation of the influence of the yaw step, at least the timing. In comparison to Figure 2.8, Figure 2.12 shows the unfiltered data as well as the filtered one for all three turbines. T1 shows between $t = 312$ s and 329 s a first reaction due to the changed wake of T0. T2 shows a first reaction between $t = 426$ s and 442 s. The filtered data shows a slightly earlier influence due to the nature of a zero phase filter.

2.A.2. AVERAGED VELOCITY IN THE NINE TURBINE CASE

Figure 2.A.2 shows the averaged wind speed in the nine turbine case from $t = 500$ s to 600 s at hub height in SOWFA. 34 slices were used to average. The wind speed is binned into 11 wind speed sections which are plotted as contours. During the time of averaging the wind direction is constant. The average wind speed of






Figure 2.13.: Averaged wind speed from $t = 500$ s to 600 s, divided into 11 speed sections.

the incoming air is at approximately 8 ms^{-1} . However, between the turbine rows, the wind speed increases to a higher level, up to 9.44 ms^{-1} in some places. This could be explained by speed-up effects: The turbines act as resistances in the flow field, and the wind speed in the place of least resistance, between the turbines, increases. The effect has been observed and described in [46] as well for instance. Due to the speed up, the turbines further downstream experience higher wind speeds than the ones in free stream and generate more energy. Figure 2.11 quantifies the effect, where T2, T4, and T5 generate significantly more energy than T6, for instance. After the wind direction change, the effect leads again T2 and T4 to generate more power. T1 is now in the situation T5 was in initially and also generates more energy. T5, however, drops to a lower level. Without an added model for effects like these, FLORIDyn (and FLORIS) will not be able to accurately describe the wind field.



3

FLORIDYN - A DYNAMIC AND fLEXIBLE FRAMEWORK FOR REAL-TIME WIND FARM CONTROL

This chapter is based on the following publication:

[47] M. Becker, D. Allaerts, and J.W. van Wingerden, *FLORIDyn - A dynamic and flexible framework for real-time wind farm control*, Journal of Physics: Conference Series (2022), 10.1088/1742-6596/2265/3/032103.

This paper presents a new framework of the FLOW Redirection and Induction Dynamics (FLORIDyn) model. It is able to dynamically simulate the wake behaviour in wind farms under heterogeneous and changing environmental conditions at a low computational cost. The novelty of this work is the improved segregation of wake dynamics and wake influence: the framework creates Observation Points (OPs) at each turbine, which propagate wind field states and turbine states downstream and follow the wind direction of the free stream velocity. These observation points cover the dynamic aspects of the simulation. The OPs, along with the stored states, are now used to derive so-called Temporary Wind Farms (TWF), which approximate the effective intra-farm wind conditions at a given location. Within these TWF, the flow conditions are homogeneous and steady state. This way, arbitrary wake models can be used to calculate the farm influence on the location. The FLORIDyn framework also provides interfaces to flow field estimators, which is tested with an effective wind speed estimator. A nine turbine case is used to highlight the quality and performance of the simulation result. Compared to its predecessor, the new FLORIDyn framework decreases the computational cost by one to two orders of magnitude, which makes it a promising candidate for real-time model predictive dynamic wind farm control.

3.1. INTRODUCTION

As wind turbines extract kinetic energy from the wind and transform it into electric energy, they leave an area of decreased wind speed in the flow field. This area is called the wake and can influence downstream turbines in a wind farm by decreasing the amount of energy these can extract from the wind. The wake is shaped by uncontrollable environmental influences and controllable turbine states. An optimization of the latter can increase the power generated by the entire wind farm: for instance an increase of the misalignment of the wind turbine with the wind direction can redirect the wake and reduce the influence on downstream turbines [48]. This control strategy is known as wake steering which is considered to be a viable solution for wind farm control [49]. One way to implement wake steering is to use a surrogate wind farm model to find the optimal control settings. The parametric FLORIS (FLOW Redirection and Induction in Steady state) model has made this approach feasible by being computationally cheap, easy to tune and implement, whilst being accurate enough to lead to performance improvements in higher fidelity simulations and wind tunnel experiments [3, 50]. A downside of FLORIS is that it neglects the dynamic behaviour of the wake and the surrounding flow and only a few implementations are capable of simulating heterogeneous environments and thus optimally suitable for dynamic wind farm control [51].

There have been a few proposals to address the lack of dynamics, one of which is the FLOW Redirection and Induction Dynamics (FLORIDyn) model, presented in 2014 by Gebraad et al. [19]. The approach is to use so called Observation Points (OPs) which travel downstream, starting at the rotor plane and inherit the turbine state and the wind field state. With this information and the boundaries of the FLORIS wake, the OPs calculate their path and propagate downstream. This way turbine state changes propagate with the OPs downstream and have a delayed effect at other turbines. While the model pioneered in its methodology and shows promising results, it also has some flaws: Firstly, the wind direction is fixed and the model does not allow for heterogeneous conditions. Secondly, the used Zone FLORIS model has since been overhauled and more capable and accurate parametric models have been developed. Thirdly, due to the OP travel behaviour, parts of the wake could overlap causing inconsistent edge cases.

In 2022, Becker et al. addressed a number of shortcomings of the Zone FLORIDyn model and extended it with new features [1]. These are, among others, the implementation of the 3D Gaussian FLORIS model, the inclusion of heterogeneous and changing flow conditions and a new method for Observation Point (OP) distribution in the wake. While the Gaussian FLORIDyn model is able to keep the computational cost low, the simulation times grow exponentially and become unfeasible for model based control approaches for a wind farm with a large number of turbines. The OPs also discretize the Gaussian shape and limit the function by their distribution. The design is quite interconnected with the Gaussian model and does not provide a simple interface to switch the parametric model as desired.

This paper aims to structurally rework aspects of the Gaussian FLORIDyn model to solve or reduce the issues it has, while maintaining its strengths. The presented work evolves a concept from the Gaussian FLORIDyn model: To simulate a wake in heterogeneous conditions, the implementation decoupled wake and flow properties in a wake and a world coordinate system. We develop this approach further to introduce Temporary Wind Farms (TWFs) which approximate

the behaviour of static turbine wakes under heterogeneous conditions. However, this formal addition has significant implications, as it changes the way FLORIDyn interacts with the underlying FLORIS model and the resource requirements of FLORIDyn. FLORIS is now treated much more as a generic model with an interface. This allows an exchange of the used wake model. This new structural framework can create a dynamic wake from any (steady state) wake representation, also from those which include secondary wake steering effects. This will also allow to include the most recent developments of the FLORIS models and such alike. The approach also lowers the computational cost in our simulations and allows a scale in which dynamic real time wind farm control with a large number of turbines becomes more viable.

The remainder of the paper is structured as follows: Section 3.2 presents the new FLORIDyn framework.

A nine turbine case is presented in Section 3.3, along with the computational cost of the framework. Section 3.4 concludes the work presented in this paper.

3.2. DEVELOPMENT OF A GENERIC FLORIDYN FRAMEWORK

First, the propagation of OPs and states is described (Sec. 3.2.1), the novel framework is then described in Section 3.2.2. How to couple the framework with measurement data is discussed in Section 3.2.3. Section 3.2.4 concludes the methodology by highlighting specific aspects of the used implementation which diverge from previous models.

3.2.1. PROPAGATION OF OBSERVATION POINTS AND STATES

This section describes the mathematics¹ of the OPs, which represent the wake of a turbine. In the new FLORIDyn framework, OPs only follow the centerline. This is in contrast to previous versions, where OPs were distributed across the entire wake area. For each turbine, a new OP is created every time step, while the oldest one is disregarded. Each OP has three sets of states: the location of the OP, denoted as \mathbf{x}_{OP} , the turbine state \mathbf{x}_T and the wind field state \mathbf{x}_{WF} . The propagation follows the concept introduced in [1] where the downwind step is calculated in the wake coordinate system \mathcal{K}_1 and then translated to the world coordinate system \mathcal{K}_0 :

$$\mathbf{x}_{OP,0}(k+1) = \mathbf{x}_{OP,0}(k) + \mathbf{R}_{01}(\mathbf{x}_{WF,\varphi}) \underbrace{[\mathbf{x}_{OP,1}(k+1) - \mathbf{x}_{OP,1}(k)]}_{\text{step in } \mathcal{K}_1}, \quad (3.1)$$

where $\mathbf{x}_{OP} = [\mathbf{x}_{OP,0}, \mathbf{x}_{OP,1}]^T$ denotes the location state of the OP with $\mathbf{x}_{OP,0} = [x_0, y_0, z_0]^T$ in the world coordinate system and $\mathbf{x}_{OP,1} = [x_1, y_1, z_1]^T$ in the wake coordinate system. Note, that x_1 denotes the downwind direction and y_1, z_1 the crosswind coordinates, y_1 from right to left, z_1 from down to up as shown in

¹The notation of Section 3.2.1 and 3.2.2 is as follows: small italic letters denote scalars (e.g. x), bold small letters denote column vectors (e.g. \mathbf{x}) and bold capital letters denote matrices (e.g. \mathbf{R}). Coordinate systems are denoted by \mathcal{K} . Square brackets organize equations or define matrices and vectors, round brackets are function inputs or properties. Lower indices of vectors first specify the parent object or type, then the coordinate system and lastly the extracted value. Depending on the context, some part of the index might be missing, but the order remains. If the extracted value is specified, the vector might reduce to a scalar and is written accordingly, e.g. $\mathbf{x}_{OP,0} \rightarrow x_{OP,0,y}$. Lower indices of matrices denote which coordinate system they transform into which, e.g. \mathbf{R}_{01} transforms \mathcal{K}_1 into \mathcal{K}_0 , such that $\mathbf{x}_0 = \mathbf{R}_{01}\mathbf{x}_1$.

Figure 3.1. The wind field state is written as \mathbf{x}_{WF} and $\mathbf{x}_{WF,\varphi}$ only refers to the wind direction φ . The rotational matrix \mathbf{R}_{01} rotates the OP propagation step in \mathcal{K}_1 around the z-axis in mathematical positive direction and therefore transforms vectors from \mathcal{K}_1 to \mathcal{K}_0 :

$$\mathbf{R}_{01}(\varphi) = \begin{bmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.2)$$

The positional update in the wake coordinates is calculated as follows:

$$\mathbf{x}_{OP,1,x}(k+1) = \mathbf{x}_{OP,1,x}(k) + \Delta t \mathbf{x}_{WF,u} \quad (3.3)$$

$$\mathbf{x}_{OP,1,y,z}(k+1) = \delta(\mathbf{x}_{OP,1,x}(k+1), \mathbf{x}_T, \mathbf{x}_{WF}) \quad (3.4)$$

where Δt is the simulation time step, $\mathbf{x}_{WF,u}$ is the free wind speed and δ denotes the deflection function. When an OP is created, $\mathbf{x}_{OP,1} = \mathbf{0}$ and $\mathbf{x}_{OP,0} = \mathbf{t}_0$ where \mathbf{t}_0 is the world location of the wind turbine rotor center.




Figure 3.1.: Visualization of the coordinate system \mathcal{K}_0 and the two \mathcal{K}_1 systems of two turbines. The figure also shows the OPs flowing downstream as well as the turbine location vectors $\mathbf{t}_{Ti,0}$ and the wind direction φ .

The description of \mathbf{x}_T and \mathbf{x}_{WF} is purposefully kept generic as these states may vary with the used parametric model. The turbine state \mathbf{x}_T can be summarized as all states that are turbine specific and are needed to calculate the wake shape. Examples are the yaw angle γ or the axial induction α . The wind field state \mathbf{x}_{WF} contains all states necessary to propagate the wake and to calculate it. In the presented formulation the free wind speed and the wind direction are mandatory, but \mathbf{x}_{WF} can also include the ambient turbulence intensity for instance. States in \mathbf{x}_{WF} are also considered to be measurable by sensors at other locations than the turbine and could be corrected by state estimation methods.

3.2.2. EXTRAPOLATION FROM OPS AND THE CREATION OF TEMPORARY WIND FARMS

The reduction of the number of OPs leads to a sparse wake description. Where other FLORIDyn formulations have the possibility to justify a nearest neighbour interpolation due to a high density of OPs, the current description cannot. Therefore

a method to extrapolate the wake influence is presented. The general problem is posed as follows:

What is the influence of the turbines \mathcal{T} at the world location \mathbf{l}_0 ?

For each turbine $T_i \in \mathcal{T}$ we choose the two OPs (OP^1 and OP^2) in front and behind the closest point on the centerline of T_i to \mathbf{l}_0 . Then we linearly interpolate between OP^1 and OP^2 to obtain OP^* in such a way that the distance to \mathbf{l}_0 is minimal:

$$w = \frac{[\mathbf{x}_{OP^2,0} - \mathbf{x}_{OP^1,0}]^T [\mathbf{l}_0 - \mathbf{x}_{OP^1,0}]}{[\mathbf{x}_{OP^2,0} - \mathbf{x}_{OP^1,0}]^T [\mathbf{x}_{OP^2,0} - \mathbf{x}_{OP^1,0}]} \quad (3.5)$$

which returns the weight $w \in [0, 1]$. The \mathcal{K}_0 location of OP^* is then given by

$$\mathbf{x}_{OP^*,0} = \mathbf{x}_{OP^1,0} + w[\mathbf{x}_{OP^2,0} - \mathbf{x}_{OP^1,0}] = (1-w)\mathbf{x}_{OP^1,0} + w\mathbf{x}_{OP^2,0}. \quad (3.6)$$

This is visualized in Figure 3.2. The weight is then also used to interpolate the other states $\mathbf{x}_{OP^*,1}$, \mathbf{x}_{T^*} and \mathbf{x}_{WF^*} , equivalent to Eq. (3.6). In the edge cases where the first or last OP of T_i is the closest OP, no interpolation is performed and $OP^* = OP^{Edge}$. This means, that if only one OP is used, the framework extrapolates from that OP and directly returns the underlying parametric wake model. After applying Eq. (3.6), every turbine in \mathcal{T} is represented by an OP^* close to \mathbf{l}_0 .

The next step is to locate \mathbf{l}_0 in \mathcal{K}_1 of turbine T_i , based on the states represented by OP^* . This is done by rotating the vector $\mathbf{l}_0 - \mathbf{x}_{OP^*,0}$ (pointing from OP^* to \mathbf{l}_0) from the world frame \mathcal{K}_0 to the wake frame \mathcal{K}_1 of T_i with the inverted rotational matrix: $\mathbf{R}_{01}^{-1} = \mathbf{R}_{01}^T = \mathbf{R}_{10}$. This vector is then added to the \mathcal{K}_1 position of OP^* :


$$\mathbf{l}_{T_i,1} = \mathbf{x}_{OP^*,1} + \mathbf{R}_{10}(\mathbf{x}_{WF^*,\varphi})[\mathbf{l}_0 - \mathbf{x}_{OP^*,0}]. \quad (3.7)$$

As a result, $\mathbf{l}_{T_i,1}$ can be calculated for all turbines $T_i \in \mathcal{T}$. The vector $\mathbf{l}_{T_i,1}$ stems from the origin of \mathcal{K}_1^* , the turbine coordinate system of T_i based on the states of OP^* . This means that if the wind field states of OP^1 and OP^2 (which OP^* is derived from) have changed since their initialization, the origins of \mathcal{K}_1^* and \mathcal{K}_1 are not at equal locations in \mathcal{K}_0 . Figure 3.2 illustrates this.

With this information we can create a Temporary Wind Farm (TWF). The TWF approximates the environment around \mathbf{l}_0 and can be seen as the effective wind farm at the requested location: This includes all turbines in \mathcal{T} as perceived by the OPs^* close to \mathbf{l}_0 . By using the data from the OPs^* , locations of the turbines in the TWF can differ from their real world locations, as shown in Figure 3.2. A new coordinate system \mathcal{K}_2 is created which is characteristically similar to \mathcal{K}_1 : The wind direction is fixed along the x_2 axis, y_2 is the crosswind direction and the z_2 axis is pointing upwards. First, we place \mathbf{l}_0 in \mathcal{K}_2 : its location can be chosen arbitrarily, but for convenience we will choose $\mathbf{l}_2 = [0, 0, l_{0,z}]^T$ which ensures that $z_2 = 0$ is ground level if elevation is not part of the simulation. The wind turbines can be located by reversing the vectors $\mathbf{l}_{T_i,1}$:

$$\mathbf{t}_{T_i,2} = \mathbf{l}_2 - \mathbf{R}_{21}(0)\mathbf{l}_{T_i,1}. \quad (3.8)$$

Since \mathcal{K}_1 and \mathcal{K}_2 share the same wind direction, $\mathbf{R}_{21}(0)$ is a 3×3 identity matrix. The turbine states are given by the respective OP^* . The wind field states however are not entirely defined, only the wind direction is fixed. In practice however, we can assume that the wind field states of the different OP^* will be very similar as they are also in local proximity to each other. In our implementation we averaged



The top figure shows the interpolation of OP^* and how it is used to calculate the position of \mathbf{l}_0 in \mathcal{K}_1^* of T_i . The wind direction is steady and \mathcal{K}_1 of the turbine is equal to \mathcal{K}_1^* of OP^* . During a wind direction change, the center line propagates to its new steady state, which leads to a change of $x_{\text{WF}^*, \varphi}$ and therefore to a temporary mismatch between the locations of the origins of \mathcal{K}_1^* and \mathcal{K}_1 in \mathcal{K}_0 (lower figure).

between the two closest OP^* , weighted by the distance to \mathbf{l}_0 . The TWF is now complete and approximates the wake and wind field conditions around \mathbf{l}_0 as a wind farm in homogeneous conditions. The TWF in \mathcal{K}_2 can be evaluated in an arbitrary wake model to return the influence of \mathcal{T} at \mathbf{l}_2 which approximates the influence of \mathcal{T} at \mathbf{l}_0 .

Note that due to the way the TWF are derived, existing steady state wake models can be used to their full extent: The presented framework does not require a certain wake shape or wake merging method. With multiple turbines in one TWF, secondary wake steering effects could be captured as well. This ensures that novel developments in these wake models can also be tested in a dynamic, heterogeneous environment.

Another aspect is that the computational load can be scaled by limiting the size of \mathcal{T} . If downstream turbines do not have an influence on upstream turbines in the used wake model, they can be disregarded from \mathcal{T} . The same goes for turbines with a significant upwind or crosswind distance to \mathbf{l}_0 . This allows a split of an entire wind farm into many smaller wind farms which can be evaluated in parallel.

3.2.3. INTERFACES AND THE IMMERSION AND INVARIANCE ESTIMATOR

The presented framework simplifies the state architecture of the simulation and purposefully treats turbines as sensors and actuators providing information to the OPs. The OPs in return provide an estimate of the intra wind farm flow. Estimators of flow field metrics can convert measurement data into metrics which the FLORIDyn framework can use and store as states in the OPs.

In this work, we present one implementation with the Immersion and Invariance (I&I) estimator as described in [52]. It returns the effective wind speed at the rotor plane $\hat{u}_{\text{eff},R}$ based on the generator torque and the rotor torque. A required component is the power coefficient table C_p which is depend on factors such as the blade pitch and tip-speed-ratio. The used look-up tables were obtained in FAST. The estimator is derived for wind turbines without yaw misalignment, which limits its use until it is adapted.

Since FLORIDyn uses the free wind speed to propagate the OPs, $\hat{u}_{\text{eff},R}$ has to be converted to $\hat{u}_{\text{free},R}$. For free stream turbines we neglect induction and blockage effects and assume that $\hat{u}_{\text{eff},R} = \hat{u}_{\text{free},R}$. The FLORIDyn framework delivers an estimate of the reduction of the wind speed for downstream turbines, which can be used to calculate $\hat{u}_{\text{free},R}$ by dividing $\hat{u}_{\text{eff},R}$ by the reduction.

3.2.4. IMPLEMENTATION

This section briefly summarizes the details of the implementation used in this work. The implementation follows in its core [1] but diverges in some aspects highlighted here.

Rotor plane discretization The rotor plane is discretized to calculate the influence of turbine wakes and sum it, weighted by the areas of the subdivided plane. We propose the use of the Isocell algorithm which splits a circle into n equally sized and regularly distributed parts [53]. The downside of this algorithm is that it can only provide sets for certain values of n . This limitation is found to be acceptable as the steps between possible values of n are relatively small ($n = 3, 12, 27, 48, 75, 108, 147, \dots$). In this work 48 rotor points (RP) were used, distributed in the y_1, z_1 plane and translated, according to the yaw orientation and turbine position, into K_0 . Only the rotor center is used as location to set up the TWF. The resulting distribution can be seen in Figure 3.1.

Thrust and power coefficient In [1] look-up tables were used to get values for the thrust and power coefficient. As these look-up tables turned out to be incomplete and needed to be limited to be useful, we decided to fall back to the actuator disc method and calculate the coefficients based on the axial induction factor.

Parameter set The presented FLORIDyn framework was also used in a parameter uncertainty quantification study [54]. The study was conducted with the FLORIS parameter set found by [51] as starting point (see Table 3.1) and is also used in this study. Additionally, the efficiency term η was set to be neutral, the yaw correction coefficient p_p was set following [1].

Added turbulence Influence of the added turbulence is implemented as presented in [24]. This formulation only defines a downwind development of the added turbulence and no crosswind part. A recommendation is to include all turbines within 2D crosswind distance and 15D upwind distance [23]. Under changing environmental conditions, this leads to sudden changes in the added turbulence levels as turbines enter or leave the proposed area.

Table 3.1.: Parameters used with the aspects they mainly influence

Wake expansion		Added turbulence			
k_a	k_b	$k_{f,a}$	$k_{f,b}$	$k_{f,c}$	$k_{f,d}$
$5.37 \cdot 10^{-1}$	$-8.48 \cdot 10^{-4}$	7.84	4.57	$4.3 \cdot 10^{-1}$	$-2.46 \cdot 10^{-1}$
Near wake length		Power			
α^*	β^*	η	ρ_p		
1.088	$2.22 \cdot 10^{-1}$	1	2.2		

Instead of looking upstream for influencing turbines, one can also look downstream to identify the turbines influenced by the wake. In the previous FLORIDyn implementation the distributed OPs were used to determine the area of influence [1]. In the here presented Framework, OPs do not cover the wake area anymore and rather extrapolate their influence from the centerline. Following this approach the added turbulence value is weighted by the Gaussian distribution of the wake. To achieve a similar width as in [1], the width factor is multiplied by three.

3.3. CASE STUDY

The presented case is a nine turbine simulation case in Section 3.3.1. The wind field is turbulent and changes the flow direction during the simulation. Section 3.3.2 discusses the computational cost aspect of the framework in comparison to the previous implementation.

3.3.1. NINE TURBINE CASE

The wind farm layout is a regularly spaced 3×3 layout in a 3×3 km domain. The turbines are placed 900 m apart from each other and 600 m from the domain edge. The validation simulation is done in SOWFA, where the DTU10MW reference turbines are simulated with the Actuator Disc Method and the domain is discretized in $10 \times 10 \times 10$ m cells without refinement areas and the time step is set to 0.5 s. The flow is turbulent with an ambient turbulence intensity of $\approx 6\%$ and a mean wind speed of 8.2 ms^{-1} . During the simulation the wind direction uniformly changes from 255° to 195° , with a constant change rate from simulation time $t = 600$ s to 900 s. The turbines remain perpendicular to the wind direction during the change. This case was also used in [1].

Figure 3.3 shows the power generated by the four downstream turbines during the simulation. The figure compares four different versions of the same simulation: The FLORIDyn framework with and without the I&I estimator, the validation data (here SOWFA) and what FLORIS would return. The grey area in the plots highlights the time during which the wind direction changes.

Looking at the difference between steady state and dynamic models, we can see that FLORIS is not able to accurately predict the timing of the wake influence of upstream turbines on downstream turbines due to the wind direction change. The power generated also shows single peaks where the validation data shows multiple occurrences of interaction. The FLORIDyn framework returns identical results during steady state. Furthermore, FLORIDyn is able to predict the timing




Figure 3.3.: Generated power by the FLORIDyn framework without and with the I&I estimator, in comparison to the validation data and FLORIS. The grey area indicates the time in which the wind direction turns. Only the four down stream turbines are plotted here, their location is marked on the left side as well as the change of wind direction.




Figure 3.4.: Flow field comparison during a wind direction change between SOWFA (1) and the new FLORIDyn framework, with I&I estimator (2) and without (3), the Gaussian FLORIDyn model (4), as well as FLORIS (5). The snapshots are taken at simulation time $t = 700$ s.

of the wake interaction. The model does lack information about the wind speed changes, which can lead to significant differences in power generated, compared to the validation data. These differences become minimal when the I&I estimator is added. Due to the dynamic nature of the I&I, the wind speed estimate is experienced with a minor phase shift and does not contain higher frequency components from the validation data. These results could potentially be even more improved by adequate tuning of the I&I parameters. The downside of the use of the I&I estimator is that it removes the FLORIDyn influence from the calculation of the power generated as it directly provides the effective wind speed. But it is in return possible to use the turbine as a sensor and to provide a wind field state estimate to the OPs.

The flow fields at hub height in Figure 3.4 compare the validation data with the FLORIDyn framework with and without I&I estimator, the Gaussian FLORIDyn

model and the FLORIS model. FLORIS shows the current wind direction, as the model immediately reacts to the new wind direction and the far wake does not slowly adapt, as is the case with the validation data and all FLORIDyn variants. The Gaussian FLORIDyn shows a mosaic-like flow field due to the nearest neighbour interpolation employed to get the influence of the OPs. The novel FLORIDyn framework allows an accurate use of the FLORIS wake, adapted to the wind direction change. Adding the I&I estimator results in a very similar wake shape but a much more detailed flow field which can incorporate locally different wind speeds based on realistic turbine data.

3.3.2. PERFORMANCE

Figure 3.5 shows the computational performance of the presented FLORIDyn framework in seven test cases, four of which have also been tested with the I&I estimator, and three which have been simulated in the previously published Gaussian FLORIDyn model [1]. The measurements tell how long it takes to simulate one time step. They are only roughly representative, as the performance will vary with the wind farm layout, the environmental conditions, the implementation and the hard- and software². The times capture only the simulation and no visualization, nor initialization. Each measurement is the mean of 250 or more simulation time steps. The 9₁ turbine case is the same case as discussed in Section 3.3.1. The 1, 2, 3 and 9₂T cases feature a turbulent wind field but steady wind direction. The 23 and 54T cases are derived from real wind farm layouts. The cases run with a constant wind speed and direction.

One observation is that the new framework, in comparison to the previous Gaussian FLORIDyn, decreases the computational cost by one to two orders of magnitude, depending on the case. The new implementation also scales better, to a point where a 54 turbine case can be simulated at a similar speed as a three turbine case in the previous implementation. This improvement allows more optimization steps in the same time, which, in return allows the application of more demanding control strategies. The TWF concept is also by design suited for parallel computation, something which was not performed in this study.

Another observation is that the added computational cost of the I&I estimator varies between the cases. This is due to the fact that the estimator runs with the frequency of the measurement data from the turbines. The 2 and 9₁ T case have 2 Hz data, the 1 T case 5 Hz and the 3 T case 25 Hz. The FLORIDyn framework runs at 0.25 Hz in all cases. Unfortunately we do not have the data to run the I&I estimator in the 9₂, 23, 54 turbine cases.

3.4. CONCLUSION

This paper presents a formal definition of the FLORIDyn framework. It decouples the wake propagation under heterogeneous, changing conditions from the wake shape description. This allows the use of wake models which are designed for steady state conditions. The new interface to a generic wake model is achieved by creating temporary wind farms which approximate intra wind farm flow at any given location. They translate the heterogeneous environmental conditions into a steady state homogeneous space in which the wake model is evaluated. This

²Simulations performed in Matlab 2021a, single threaded on a laptop (Intel Core i7 vPro 10th Gen, 16GB RAM).




Figure 3.5.: Computational performance of the new FLORIDyn framework with and without the I&I estimator in comparison to the Gaussian FLORIDyn model. Note the logarithmic time scale and the number of turbines for the different cases. The FLORIDyn framework values are mean value of 10 consecutive simulations.

allows the dynamic use of most recent steady state wake model implementations with minimal effort and makes the concept usable for future generations of engineering wake models.

This paper also demonstrates how measurements from a turbine can be included into the FLORIDyn framework: The I&I effective wind speed estimator is used to convert the turbine into a sensor for the simulation, which significantly decreases the differences in the power generated and improves the intra wind farm flow estimate.

Compared to the previous Gaussian FLORIDyn model for heterogeneous conditions, the FLORIDyn framework is able to decreased computational effort by one to two orders of magnitude. This is a necessary step in order to provide real time control inputs in a realistic wind farm scenario with a large number of turbines. The temporary wind farms also provide a basis for parallel computing, something which was not part of this research.

For future work we propose to implement the framework into a model based closed loop control strategy. Within the FLORIDyn framework a suitable engineering model can be used to evaluate desired quantities of interest under changing environmental conditions. This could be for instance a wake steering strategy to maximize power while taking loads into account. Additionally other flow field estimators can also be implemented to improve the connection of FLORIDyn with data from wind farms.



4

A DYNAMIC OPEN-SOURCE MODEL TO INVESTIGATE WAKE DYNAMICS IN RESPONSE TO WIND FARM FLOW CONTROL STRATEGIES

This chapter is based on the following publication:

[55] M. Becker, M. Lejeune, P. Chatelain, D. Allaerts, R. Mudafort, and J.W. van Wingerden, *A dynamic open-source model to investigate wake dynamics in response to wind farm flow control strategies*, Wind Energy Science (2025), 10.5194/wes-10-1055-2025.

Wind farm flow control (WFFC) is the discipline of manipulating the flow between wind turbines to achieve a farm-wide goal, like power tracking, load mitigation, or power maximization. Specifically, steady-state control approaches have shown promising results in both theory and practice for power maximization. But how are they expected to perform in a dynamically changing environment? This paper presents an open-source wake modeling framework called OFF. It allows the approximation of the performance of WFFC strategies in response to environmental changes at a low computational cost. It is rooted in previously published dynamic parametric engineering models and offers a flexible and adaptable platform to explore these models further. The presented study tests the modeling framework by investigating the performance of different wake steering controllers in a 10-turbine wind farm case study based on a subset of the Dutch wind farm Hollandse Kust Noord (HKN). The case study uses a 24-hour wind direction time series based on field data and verifies subsets of the time series in LES. The results highlight how dependent yaw travel is on the controller settings and suggest where users can strike a balance between power gains and actuator usage. They also show the structural differences and similarities between steady-state and dynamic engineering models. The comparison to LES shows what time scales the surrogate models cover and how accurately. While steady-state models capture turbine power signal dynamics up to $\approx 1/570$ Hz, the dynamic wake description can predict dynamics up to $\approx 1/360$ Hz with a better correlation and normalized root-mean-square-error. Further results show that the dynamic wake description is mainly advantageous over steady-state wake models for shorter periods (< 20 min). The paper also opens up the discussion about the effectiveness of wind farm flow control in a time-marching manner as opposed to a steady-state viewpoint.

4.1. INTRODUCTION

Wind energy is an essential part of the modern renewable energy mix and, therefore, part of the increasing share of energy that is covered by renewables. With this increasing share comes a higher responsibility. Where previously only individual turbines would contribute to the electrical grid, numerous wind farms provided 19% of the electricity demand in the EU in 2023 [56]. With this increased relevancy, the question arises of whether wind farms are used to their full extent. Their efficiency could be limited, among other reasons, by unintended turbine downtime, maintenance or non-ideal operation. Wake losses are included in the latter, as front-row turbines extract kinetic energy from the wind, and they inevitably slow down the flow behind them. The turbines downstream thereby experience a lower wind speed and generate less power in response. To combat this effect, wind farm flow control (WFFC) methods focus on lessening the losses induced by wakes. This is achieved by modifying the behavior of the turbines from a greedy control approach to a collaborative one.

Multiple control approaches exist to address this issue. They can be sorted by the degrees of freedom they use: (i) the blade pitch (e.g. ., [57], [58]) (ii) the generator torque (e.g. [59]) and (iii) the (mis-)alignment of the turbine with the flow (e.g. [60], or [61]). Broadly speaking, (i) and (ii) change how much energy is extracted from the flow field. Applied dynamically, the blade pitch can also increase wake mixing behind the turbine, which leads to a faster wake recovery. In contrast, using (iii), the alignment of the rotor allows the controller to deflect the wake in the lateral direction. This control strategy can be used to direct the wake away from downstream turbines and is referred to as wake steering. The remainder of the paper focuses on this effect and methods to determine the effectiveness of control strategies using wake steering.

To research, test and optimize control strategies for wind farms, surrogates of the real plant are needed. This mitigates risks, lowers costs, increases flexibility and makes the problem more accessible. Alongside wind tunnel experiments (e.g. , [16, 62]), simulations are the predominant way to approximate wind farm behavior. Within the world of simulations, three groups can be distinguished: high-, medium-, and low-fidelity simulations. High-fidelity models, such as large-eddy simulation (LES), provide the most accurate approximation of the flow field (e.g. [20, 63]). This does come at an increased computational cost, which has confined their application to the verification or exploration of new phenomena not yet captured by lower-fidelity models. At the other end of the spectrum, low-fidelity simulations reduce the wake behavior to a set of simple analytical equations that are efficient to solve. This, however, means that they can only describe what they have been designed for: typically a single time-averaged snapshot of the flow field (e.g. [2, 16]). Low-fidelity models are therefore routinely used to, for instance, optimize the orientation of all turbines in a wind farm for the entire wind rose, to make estimates of the annual energy produced (AEP) or to optimize the wind farm layout.

Growing concerns about fatigue effects on wind turbine integrity, along with the rising need for ancillary service provision, have driven recent research toward a new generation of dynamic medium-fidelity models. These models are designed to address more immediate and transient phenomena, effectively bridging the gap between high- and low-fidelity approaches. By capturing the critical dynamics of high-fidelity simulations at a fraction of the computational cost, they move beyond

steady-state assumptions, unlocking new possibilities for wind farm operations. Key applications include, for example, intra-hour power production predictions for grid regulation (e.g. [64]), as well as multi-objective wake steering strategies that optimize the power output while simultaneously mitigating the turbine's loads (e.g. [65]).

Medium-fidelity wake models are primarily categorized by the equations they use to model flow physics, balancing computational cost with accuracy. While 2D linearized Reynolds-averaged Navier-Stokes (RANS) methods have demonstrated some initial success at estimating simple wake states, they have been shown to improperly account for wake deflection [66]. In contrast, free-vortex methods (e.g. [67], [68], or [69]) explicitly resolve vortex dynamics, providing deeper insights into large-scale wake behavior. This capacity, to account for phenomena such as wake deflection and wake curling, makes free-vortex methods ideal candidates to investigate wake steering. However, the computational burden associated with these methods makes them unsuitable for large parameter spaces, such as those encountered in offshore wind farms involving dozens of turbines. Additionally, they tend to become numerically unstable for large distances and are, therefore, limited in terms of the wake length they can describe accurately.

The dynamic wake meandering model, initially proposed by [17], also opts for a Lagrangian parametrization of the wake, describing it as a cascade of velocity deficits without explicitly solving vortex dynamics. Since its introduction, the DWM approach has been further calibrated and validated by numerous studies comparing it against both numerical and field data [18, 70, 71]. Building on these early successes, it has been integrated into simulation software such as FAST-Farm [72] and HAWC2FARM [73]. More recently, the DWM model has been reinterpreted into a series of lighter, control-oriented wake modeling frameworks that include FLORIDyn [1, 19, 47, 74], OnWARDS [75], UFloris [76], and SWiPLab-WFM [77]. A common feature of these models is that they all adopt a Lagrangian description of the flow while relying on engineering wake models to capture the wake's influence. However, though similar, these models take different paths notably regarding how they handle the ambient flow field and wake deflection. They also differ in terms of the steady-state surrogate wake model, which is generally fixed for the presented designs. And, while steady-state models have been summarized in unified toolboxes (FLORIS [78], PyWake [79], or FOXES [80]), dynamic engineering models have not.

The purpose of OFF (abbreviation based on OnWARDS, FLORIDyn and FLORIS), the dynamic wake modeling framework presented in this paper, is to provide a unified, open-source toolbox that allows for easy comparison between different implementations. Specifically, the framework aims to:

- design and implement an interface with established steady-state models, such as FLORIS [78] or PyWake [79];
- provide a framework for prototyping Lagrangian dynamic wake models through standardized input-output structures, facilitating the replicability of results;
- offer accessibility through open-source code written in Python.

Such a tool shall eventually allow for benchmarking and comparisons of dynamic and steady-state wake model designs and for further exploration and development

of dynamic WFFC strategies at a low computational cost (as already utilized by, e.g. [81, 82]). Further scientific contributions of this paper are:

- an investigation into the timescales captured by steady-state wake models versus those captured by dynamic wake models, providing insights to help users make informed choices based on their specific needs;
- a verification of the presented code using LES in a neutral ABL with a 10-turbine wind farm;
- a dataset based on a total of 54 h of LES simulation with varying controller settings and changing wind directions to use for further wake model analysis and synthesis.

The following paper is split into five sections. While Section 4.1 introduces the context of the work, Section 4.2 describes the presented model and its architecture, as well as details of the implementation used to generate the results from this paper. Section 4.3 then presents a case study where a selection of yaw steering controllers are investigated in the presented model, followed by Section 4.4, where a selected range of controllers are implemented in the LES. The section goes on to compare the LES results to the results predicted by the dynamic model as well as by the steady-state model. Lastly, Section 4.5 concludes the paper and suggests pointers for future work.

4.2. MODEL DESCRIPTION

The framework called OFF is designed to run generic particle-based dynamic wind farm flow simulations using three sets of states: (i) turbine states \mathbf{x}_T , (ii) ambient states \mathbf{x}_{amb} and (iii) observation point (OP) states \mathbf{x}_{OP} . Turbine states consist of all states necessary to describe the turbine's impact on the wake, e.g. the turbine yaw angle and its axial induction. The ambient states characterize the flow field, with information about wind speed, direction, and ambient turbulence intensity. The observation point states finally map the world (*i.e.* inertial) coordinate system to the wake one, thereby allowing the reconstruction of a snapshot of the flow field across the wind farm. The states are then updated through three consecutive steps - prediction (Equation(4.1)), correction (Equation (4.2)), and control (Equation (4.3)):

$$[\mathbf{x}_T(k), \mathbf{x}_{amb}(k), \mathbf{x}_{OP}(k)] = f_{\text{prediction}}(\mathbf{x}_T(k-1), \mathbf{x}_{amb}(k-1), \mathbf{x}_{OP}(k-1), \mathbf{c}), \quad (4.1)$$

$$[\mathbf{x}_T(k), \mathbf{x}_{amb}(k), \mathbf{x}_{OP}(k)] = f_{\text{correction}}(\mathbf{x}_T(k), \mathbf{x}_{amb}(k), \mathbf{x}_{OP}(k), \mathbf{m}(k), \mathbf{c}), \quad (4.2)$$

$$\mathbf{x}_T(k) = f_{\text{control}}(\mathbf{x}_T(k), \mathbf{x}_{amb}(k), \mathbf{x}_{OP}(k), \mathbf{m}(k), \mathbf{c}), \quad (4.3)$$

where \mathbf{c} denotes a set of parameters, k the time step, and \mathbf{m} a set of measurements. The prediction step advances the model by itself: it propagates and updates the information gathered at the previous time steps. The correction step then uses the current measurements to alter the predicted states, partially reconciling them with the real-flow field. The last step finally determines the control actions the turbine takes based on the current state and measurements.

Summarizing, the OFF framework offers a prototyping environment for the development and assessment of new dynamic flow modeling strategies. The




Figure 4.1.: Nested software architecture used for the results presented in this paper: the OFF framework provides the interface to the wake solvers, as well as the controller. In this paper, the FLORIDyn framework is used to model the state dynamics, like the wake advection. The framework approximates the flow field at the location of each turbine and uses FLORIS to calculate measurements like effective wind speeds and power generated.

4

update steps are kept generic, thereby allowing the user to specify its own update strategy, for instance, by switching the dynamic solver or wake model used. Figure 4.1 depicts the version of the code used here that follows the FLORIDyn framework and uses FLORIS v4 as a surrogate model. The implemented update steps are further detailed in the following sections: Section 4.2.1 further specifies the FLORIS and FLORIDyn models used, and Section 4.2.2 explains how external data are fed into the simulation. Lastly, Section 4.2.3 introduces the control law used in this paper.

4.2.1. PREDICTION: WAKE AND TURBINE MODELING

The prediction step is segmented into three parts: (i) propagate the states, (ii) run the steady-state surrogate model to get turbine measurement predictions and OP advection speeds for the next time step, and (iii) retrieve information relevant to the controller. The states related to a single turbine T at the x, y, z location $l_{T,x}, l_{T,y}, l_{T,z}$ are propagated as follows:

$$\mathbf{x}_T(k) = \mathbf{A}_1 \mathbf{x}_T(k-1), \quad (4.4)$$

$$\mathbf{x}_{\text{amb}}(k) = \mathbf{A}_1 \mathbf{x}_{\text{amb}}(k-1), \quad (4.5)$$

$$\mathbf{x}_{\text{OP},x}(k) = \mathbf{A}_2 [\mathbf{x}_{\text{OP},x}(k-1) + \Delta t \mathbf{x}_{\text{amb},u}(k-1)] + [l_{T,x}, 0, \dots, 0]^T,$$

$$\mathbf{x}_{\text{OP},y}(k) = \mathbf{A}_2 [\mathbf{x}_{\text{OP},y}(k-1) + \Delta t \mathbf{x}_{\text{amb},v}(k-1)] + [l_{T,y}, 0, \dots, 0]^T,$$

$$\mathbf{x}_{\text{OP},z}(k) = \mathbf{A}_2 \mathbf{x}_{\text{OP},z}(k-1) + [l_{T,z}, 0, \dots, 0]^T, \quad (4.6)$$

$$\mathbf{A}_1 = \begin{bmatrix} 1 & 0 & & \mathbf{0} \\ 1 & 0 & & \\ & \ddots & \ddots & \\ \mathbf{0} & & 1 & 0 \end{bmatrix}, \quad \mathbf{A}_2 = \begin{bmatrix} 0 & 0 & & \mathbf{0} \\ 1 & 0 & & \\ & \ddots & \ddots & \\ \mathbf{0} & & 1 & 0 \end{bmatrix}, \quad (4.7)$$

where the matrices \mathbf{A}_1 and \mathbf{A}_2 handle the state propagation. With \mathbf{A}_1 , all states besides the first one are propagated one entry further, and the last one is disregarded. The state closest to the turbine is effectively doubled. With \mathbf{A}_2 the first state is not doubled but overwritten by a new input. States propagated with \mathbf{A}_1 do not have a new input yet; e.g., there is no new wind speed value available at this time in the simulation cycle. Therefore, the current wind speed is kept as a prediction. The OP position states, however, do have a new input, which is the

rotor center location, why they are propagated with \mathbf{A}_2 . Equation (4.6) updates them with the turbine location $l_{T,x}$, $l_{T,y}$, $l_{T,z}$, referring to the rotor center, as a new state. In a floating-turbine scenario, this could be used to induce a changing turbine and wake location due to repositioning. Note that similar, more detailed state-space descriptions can be found in [19, 76, 83]. A difference between these formulations and the one employed in OFF is that OFF's formulation does not include vertical or horizontal OP deflection based on the yaw and tilt angle of the turbine. Rather, the impact of yaw and tilt turbine misalignment on the wake shape is solely simulated in the wake model. The code internally decomposes the wind speed and direction into its u and v components to avoid unexpected behavior when switching between 360 and 0 deg. These are then used along with the time step Δt to advance the location of the OPs through a Lagrangian update; see Equation (4.6). The w component is ignored for simplicity. Accounting for the vertical deflection of the wake center might become necessary in some contexts, e.g., for simulations including terrain. However, it was not deemed necessary for the application presented here, i.e., an offshore wind farm with no tilting. Note that this implementation also assumes that the OP advection speed is equal to the freestream wind speed. Alternatives are the introduction of a constant fraction of the wind speed, see for instance [84] or the use of the effective wind speed predicted by the wake model, see for instance [85]. One may also decide to decouple ambient particle advection from the OP advection, thereby allowing the capture of additional wake dynamics such as wake meandering [75]. These approaches, however, increase the computational cost of the model, as it requires the evaluation of the wake equations for every OP at every time step. Equation (4.5) does not include inputs as new ambient state information is introduced via the correction step; see Section 4.2.2. Similarly, new turbine states may be introduced in the correction or in the control step; see Section 4.2.3.

After the states are propagated, the wake model is evaluated to retrieve predicted measurements. This process uses the so-called TWF, which provides a localized approximation of the ambient and wake conditions at a specific turbine location. More specifically, the temporary wind farm (TWF) maps the current dynamic state of the simulation to the corresponding steady-state configuration at any desired position, making it interpretable by the underlying wake model, i.e. FLORIS. For more details, we refer to [47]. A block diagram example is given in Figure 4.2. The graph shows the equivalent of a three-turbine wind farm where turbines T1 and T2 wake turbine T3. Turbines T1 and T2 both receive input from the wind field; add their own states; and pass them on to the first OP, which adds its own states. The set of the three state vectors is then propagated downstream. Downstream, T3 is subject to the wakes of T1 and T2. To calculate the wind speed reduction, one ghost OP is interpolated for each impacting wake. The ghost OP is based on the two closest OPs in the wake and minimizes the distance between the chain of OPs and the turbine T3. Its state is a distance-based interpolation of the two parent OPs. The state information of the ghost OPs subsequently approximates the ambient conditions and wind farm surrounding turbine T3. The TWF is then passed on to the steady-state surrogate model for evaluation. This returns predicted measurements like the effective wind speed and power generated. At each time step, a new individual TWF is generated for each of the n_T turbines. This leads to n_T TWF simulations, where each of them contains n_T turbines. The resulting computational cost is discussed in Section 4.4.6. This work interfaces with the FLORIS toolbox and uses the Gauss Curl Hybrid model [86] with default




Figure 4.2.: Schematic of the state transportation of turbine states, ambient states, and observation point states in a three turbine example. T1 and T2 wake T3. The OPs closest to T3 in the wakes of T1 and T2 are used to create a temporary wind farm (TWF) to simulate the resulting conditions for T3 in the wake model. The colored cubes indicate the states that are passed between the different elements of the software.

settings and parameters. No parameter tuning was performed to represent the performance achievable with the default settings. The turbine model within FLORIS is based on the $c_p(u)$ and $c_t(u)$ tables (u being the wind speed ahead) of the DTU 10 MW [33], corrected with the blade element momentum theory based cosine-loss law for yaw misalignment. [87, 88]. Specifically, the classical value of 1.88 is retained for the cosine power-loss law exponent. We nonetheless acknowledge that this constant power-loss model does not account for the variability of operating conditions and will therefore likely affect the optimal steering angles computed, as noted by [89].

4.2.2. CORRECTION: LINKING MEASUREMENTS AND STATES

In this work, only ambient states are corrected. Schemes to correct the wake location exist [90, 91] but are outside of the scope of this paper. Three ambient states are considered in the presented version of the model: wind direction, wind speed, and ambient turbulence intensity. Out of these three, only the wind direction varies in the presented simulations. By design, OFF assumes that measurements are taken at the locations of the turbines. The correction step has to alter the simulation states \mathbf{x}_{amb} to incorporate the new information provided. The basic assumption is made that the wind direction changes uniformly for the entire wind farm. As a result, all wind direction states are overwritten with the new measurement, which is assumed to be noise-free. Practically, this is due to the fact that the measurements used for the wind direction in the experiments stem from a single location; more details are given in Section 4.3.1. In an alternative setup with more measurement locations available, a sensor fusion strategy is necessary. Possible approaches to use turbine measurements to correct ambient states in the field exist, like a weighted map, as done by [23], a Kalman filter, as done by [21]; or an Ensemble Kalman filter, as applied by [83].

4.2.3. CONTROL: STATE-BASED DECISION MAKING

The employed controller is based on [92] and implements a yaw steering dead-band controller that relies on a look-up table (LuT) aggregated using FLORIS. Specifically, this LuT associates each wind direction with a set of optimal steering angles. In a dynamic environment, the controller now has to apply the optimized angles based on the current (estimated) ambient conditions. To this end, the controller has an estimate of the wind direction $\hat{\phi}$, which is updated based on its own value in comparison with the measured wind direction. The estimated wind direction is then used to evaluate the LuT and provide new set points. More precisely, the yaw steering control law is formulated as follows:

$$\varphi_f(k) = f_{\text{filt}}(\varphi_m(k), \varphi_m(k-1), \dots, \varphi_m(0)) \quad (4.8)$$

$$\hat{\phi}(k) = \begin{cases} \varphi_f(k) & \text{if } |\varphi_f(k) - \hat{\phi}(k)| > \varphi_{\text{lim}} \text{ or } k_i \left| \sum_{l=k-1}^{\tau} \varphi_f(l) - \hat{\phi}(k) \right| > \varphi_{\text{lim}} \\ \hat{\phi}(k-1) & \text{otherwise} \end{cases} \quad (4.9)$$

$$\gamma(k) = f_{\text{LuT}}(\hat{\phi}(k)), \quad (4.10)$$

where τ marks the time step of the last update of $\hat{\phi}(k)$ to a new value. The measured wind direction at the time step k and its filtered version are denoted by $\varphi_m(k)$ and $\varphi_f(k)$, respectively. The control law has four elements that need to be supplied: the low-pass filter, f_{filt} ; the dead-band width, φ_{lim} ; the integration coefficient k_i ; and the LuT, f_{LuT} . These elements determine the behavior of the wind farm, and their adequate tuning is a prerequisite to efficient wake steering. The selection of the parameters φ_{lim} and k_i is the subject of the case study presented in Section 4.3. The f_{filt} function is omitted for simplicity; instead we assume an ideal noise-free measurement of the wind direction. The LuT is first populated using the serial-refine yaw optimizer integrated into FLORIS [93]. While the presented control law focuses on wind direction changes, for completeness, the provided LuT also includes inputs for other freestream atmospheric conditions, such as hub-height turbulence intensity (TI) and free wind speed. These parameters are kept constant in the case study discussed in Section 3. During the LuT

creation, TI is kept constant at 6 %, the wind direction is discretized into 1 deg bins, and the wind speed from 6 ms^{-1} to 10 ms^{-1} in 1 ms^{-1} steps. The baseline controller follows the same update law, with the difference that it enforces turbine alignment with $\hat{\phi}(k)$. The controllers are continuously updated with every 5 s time step of the simulation; the limits of the intentional misalignment with the main wind direction are set to ± 30 deg.

4.3. CASE STUDY

The “Case study” section is split into two parts: Section 4.3.1 discusses the selection and processing of the field data and the resulting simulation conditions. Section 4.3.2 then showcases the use of the OFF model to predict the performance of controllers and how a pre-selection can be made from a large number of controllers.

4

4.3.1. SIMULATION SETUP




Figure 4.3.: (a) The full 23 h and 45 min wind direction time series investigated in this work. The series is based on field data recorded by a vertical LiDAR at the HKN site during the 28th of March 2023 [94], depicted in grey. The low-pass filtered data is given in black. Three marked subsets of the time series have been simulated in LES for verification purposes. Each LES TF has a length of 3 h, along with a 20 min initialization period. Critical wind directions are marked in (a) and depicted relative to the farm layout in (b).

The case study is based on the southwestern corner of the Hollandse Kust Noord (HKN) wind farm, which consists of 10 turbines, here modeled as DTU 10 MW reference turbines with a diameter of $D = 178.3 \text{ m}$ [33]. The layout has been scaled to preserve the same relative distances between the turbines compared to the original ones. It features three critical wind directions for which three or more turbines stand in line, namely for $\phi \approx 175$, 201, and 265 deg. To effectively challenge the controllers, a wind direction time series that is both realistic and includes variations across all three directions (along with smooth transitions between them) is desirable. Accordingly, to drive the simulation, we use 23 h and 45 min of data recorded by a vertical ZephIR 300M wind lidar at the HKN site on March 28, 2023, as shown in Figure 4.3 [94]. This date is before the wind farm went online, which happened in December 2023¹. The lidar provides horizontal and vertical wind speeds, along with wind directions, at various heights.

¹www.crosswindhkn.nl, accessed 28th of October




Figure 4.4.: Lidar location within the HKN wind farm site with respect to the neighboring wind farms *Prinses Amalia Windpark* (PAW) and *Egmond aan Zee* (EaZ), as well as its distance to the closest considered turbine. The measurements used in this study range from 172 deg to 304 deg, part of which, 190 to 211 deg, may be influenced by PAW. Note that the data used were recorded before HKN went online.

For this study, measurements at 108 m and 133 m were used to interpolate the wind direction at a hub height of 119 m. In order to recover the underlying wind direction changes, the ensuing signal was then zero-phase low-pass-filtered using a fourth-order Butterworth filter with a cut-off frequency of 1/600 Hz, equivalent to [95]. The filtered output, as well as the wind direction input for the precursor, was eventually fed to the yaw steering controllers. For the controller, this results in an unrealistic noise-free signal, which would otherwise be a function of a filter or distributed estimation algorithm, e.g., [96–98]. Since this work aims to demonstrate the surrogate model capabilities and not necessarily the effectiveness of an integrated wake steering controller, the added complexity of a wind direction estimator has been left out. Figure 4.4 depicts the lidar location in the context of the HKN wind farm site and its closest neighboring wind farms². The figure shows that the used wind direction range overlaps with the direction in which the Prinses Amalia Windpark is located, which may have an impact on the measurements. Therefore, for the purposes of this paper, changes in wind speed are neglected, and a constant mean wind speed of 8 ms^{-1} is imposed for all simulations. This wind speed corresponds to the turbine's below-rated operation region, where the impact of wake losses is most significant, thereby offering the greatest potential for power maximization using wake steering. The OFF simulations ran with a shear coefficient of 0.12, a turbulence intensity of 6 % and no veer. Each turbine uses 200 OPs to describe the wake. With a time step of 5 s and a freestream wind speed of 8 ms^{-1} , this results in 8 km of simulated wake, or 44.9 D, which reaches beyond the boundaries of the simulated farm (approximately $5 \times 5 \text{ km}$ region).




Figure 4.5.: Comparison of the trigger condition that leads to an updated wind direction based on Equation (4.9). Red means that the controller is updated more often based on an exceeded dead band, and blue means that the integrated error crosses the threshold more often. Marked squares indicate controller settings selected for verification in Section 4.4

4.3.2. PREDICTED CONTROLLER PERFORMANCE

The controller Equation (4.9) updates the wind direction estimate based on either of two conditions: (i) the difference between the current wind direction and the measured direction is larger than ϕ_{lim} or (ii) the integrated error exceeds the threshold. To ensure a sensible range of parameters, we investigate the balance between these two conditions: Fig. 4.5 compares which of the two triggers dominates and causes a LuT reevaluation. The results show that the chosen range of $\phi_{lim} \in [2, 10]$ and $k_i \in [0.01, 0.09]$ leads to both cases: either a predominant role of the threshold or one of the integration constant.

The selected ranges of ϕ_{lim} and k_i with a 1 deg and 0.01 discretization, respectively, lead to 81 possible combinations of dead-band settings for two types of controllers, LuT and Baseline (BL). All 162 controllers are evaluated using OFF with the results reported in Figure 4.6. The figure displays the controller performance in three dimensions: (i) energy generated, (ii) number of yaw actuator activations, and (iii) accumulated yaw travel. Figure 4.6 (a) compares the activations with the energy generated, (b) the energy with the yaw travel, and (c) the yaw travel with the activations. All three figures are colored based on their ϕ_{lim} setting. Looking at the baseline controllers in Figure 4.6 (a), it becomes apparent that a smaller ϕ_{lim} results in many more activations but not in an increase in energy. This is a result of a power curve that has little sensitivity to small yaw angle misalignment, possibly highlighting the need for more adequate power-loss exponent parametrization [89]. On the other hand, the LuT-controlled cases still benefit

²Adapted from map.4coffshore.com/offshorewind/, accessed 28th of October




Figure 4.6.: (a-c) Unfolded three-dimensional performance comparison of the dead-band controllers across the full simulated time frame in OFF. Next to the energy generated by the 10-turbine wind farm, there is the accumulated yaw travel in degrees and the number of times the yaw actuators are activated. The baseline controllers are colored in different shades of red, based on ϕ_{lim} . The LuT controllers are colored in blue, respectively.

from the increased number of activations, but with diminishing returns. Notably, there is little difference in the number of activations between baseline and LuT controllers. This is due to the fact that Equation (4.9) updates the wind direction estimates for baseline and LuT controllers alike. In contrast, the LuT controllers accumulate a much larger amount of yaw travel than the baseline cases, as depicted in Figure 4.6 (b). This is to be expected as the baseline controllers only drive the turbines to full alignment, while the LuT may vary between large positive and negative misalignment angles. Figure 4.6 (c) shows the relation between activations and yaw travel. The plot completes the picture drawn by (a) and (b): while the number of actuator activations may be similar between baseline and LuT controllers, the yaw travel is not. From these results, one could start to deduce which controllers fall within a reasonable range for set turbine limitations.

For instance, if there is an average yaw activation budget of 10 times per hour per turbine, the number of relevant controllers can be reduced. In this case, $23.75 \text{ h} \cdot 10 \text{ turbines} \cdot 10 \text{ activations per turbine}$ leads to a maximum of $= 2375$ activations, which limits the dead-band width at $\varphi_{\text{lim}} \geq 5$ deg. The results show that if yaw travel and turbine misalignment are not of concern, a LuT controller may result in a significant improvement in energy generated.




Figure 4.7.: LuT controller performance normalized by the respective baseline controller with identical φ_{lim} and k_i settings. Three marked settings along the min-max approximated Pareto front are chosen for verification. The coloring is based on φ_{lim} .

In this work, we select the controllers for verification based on the performance difference due to the switch from Baseline to LuT control. Figure 4.7 shows how the addition of wake steering, while maintaining the same φ_{lim} and k_i , increases the amount of yaw steering in comparison to the increase in farm energy. The minimize-yaw-travel and maximize-energy approximated Pareto front indicates several candidates that offer a trade-off between the increase in energy and the resulting increase in yaw travel. Three combinations of φ_{lim} and k_i along the front are selected for LES verification: one that yields a steep increase in energy for a relatively low increase in yaw travel ($\varphi_{\text{lim}} = 10$ deg and $k_i = 0.05$), one that tries to achieve the maximum energy possible ($\varphi_{\text{lim}} = 2$ deg and $k_i = 0.09$), and one intermediate configuration ($\varphi_{\text{lim}} = 5$ deg and $k_i = 0.02$).

Next to the results presented in Figure 4.6 and 4.7, which summarize the overall performance, a wind direction resolved investigation of the results can also be useful. Figure 4.8 (a) shows the energy generated by the baseline and LuT controllers with $\varphi_{\text{lim}} = 5$ deg, and $k_i = 0.02$ versus the wind direction. More specifically, a sliding time window of 600 s is used to calculate the energy, as well as the mean wind direction and wind direction change. The result is a smooth transition between multiple 10-minute average bins. The energy data are plotted over the mean wind direction and, therefore, go back and forth along the x axis (compare Figure 4.3). In direct comparison, it is evident that the LuT manages to outperform the baseline controller as expected for large parts of the wind direction,




Figure 4.8.: (a) Energy generated by the wind farm, calculated based on the power integrated over a sliding time window of 600 s. The energy is plotted over the mean wind direction $\bar{\varphi}$ during the 600 s for both LuT and BL control. The resulting wind farm efficiency is given in (b) and (c). Next to the wind farm efficiency, (b) also depicts the predicted LuT steady-state wind farm efficiency. In (c), the efficiency is given as color, while the y axis denotes the mean wind direction change $\dot{\varphi}$ over 600 s. The controller settings are $\varphi_{\text{lim}} = 5$ deg, and $k_i = 0.02$.

though, not for all of them. Figure 4.8 (b) depicts the wind farm efficiency as the ratio of the energy generated by the LuT divided by the baseline energy. The data show that the LuT-driven controller shows advantageous behavior for wind directions between 160 deg and 220 deg but struggles to consistently outperform the baseline in the wind direction transitions between 220 deg and 300 deg. Figure 4.8 (b) also depicts the wind farm efficiency as predicted by FLORIS during the LuT creation, so under ideal steady-state conditions. The difference between the achieved wind farm efficiency and the predicted one shows that the changing turbine states and wind direction state can lead to suboptimal performance and that the wind farm efficiency predicted by the LuT is, in most cases, an upper limit, only achievable under steady-state conditions [99]. The occasional localized overshoots beyond this performance envelope can be attributed to the dynamic nature of the simulations. For instance, in the absence of wake steering, a downstream turbine aligned with the wind direction would always operate within the wake of the upstream turbine. However, in a dynamic setup, transient wind direction changes may temporarily shift the wake, allowing the downstream turbine to operate under improved conditions and produce more power than in the steady-state scenario. Nevertheless, these overshoots are temporary, eventually converging back to the steady-state value or lower. This observation highlights the need for dynamic wake models to optimize wind farm control strategies during transient periods. Lastly, Figure 4.8 (c) shows the wind farm efficiency over the mean wind direction, as well as the mean wind direction change. This serves as an approximated state-space representation of the wind

direction and how it influences the wind farm performance. Since the y-axis depicts the wind direction change, the state of the wind direction moves left in the lower half of the plot and right in the upper half. In conclusion, the performance of a wake steering controller is not trivial to assess in a time-marching simulation due to changes in the flow field and in the turbine state. As a result, the wind farm can exhibit very different performance for the same wind direction and wind speed.

4.4. HKN CASES THAT UNDERWENT LES

This section verifies the selected controllers from Section 4.3.2 across the three subsets of the 24-hour period simulated in OFF and FLORIS. The OFF results are compared to both the LES and FLORIS, allowing the effect of the added dynamics to be investigated. Section 4.4.1 further introduces the LES setup and the three time frames. Section 4.4.2 - 4.4.4 investigate the power generated on a turbine, farm and statistical level respectively. This is followed by Section 4.4.5, where the energy generated is compared between the simulations.

4

4.4.1. LARGE-EDDY SIMULATION

The 10-turbine wind farm is simulated as actuator discs in a $5 \times 5 \times 1$ km simulation domain in SOWFA [20]. The domain is discretized in $300 \times 300 \times 100$ cells and simulated with a time step of 0.5 s. A grid resolution of $16.6 \times 16.6 \times 10$ m was chosen to balance computational cost and accuracy. Given the turbine rotor diameter of 178.3 m, this results in a normalized cell width of $\Delta x = \Delta y = 0.094 D$, but since the turbines are often diagonally oriented in the domain during the simulation, there is a worst-case ratio of $\sqrt{2}\Delta x = 0.132 D$. The neutral turbulent precursor is developed over $3 \cdot 10^4$ s. A surface roughness of 0.0002 m enforces a horizontal turbulence intensity of $\approx 6.2\%$ at hub height. The initial wind direction is kept constant at 225 deg during the precursor to allow changes of ± 45 deg in the successor phase, using the same southern and western inflow planes. Three 3 h successor phases underwent LES, as marked in Figure 4.3. A 1200 s spin-up phase with fixed wind direction is first run in order to fully propagate the wake, after which 10800 s of the low-pass-filtered field data is used to uniformly change the wind direction. All three time series are offset to start with 225 deg, while the wind farm layout is rotated in the LES thereby ensuring the same precursor can be used across all three simulations. The veer of the precursor is < 2 deg across the rotor plane, and the shear exponent is ≈ 0.075 . Figure 4.9 shows the wind farm in the rotated domain and a qualitative visualization of the wind directions during the simulation. The latter is achieved by a pizza-shaped histogram with bins of 2.5 deg in width, translated onto the position of each turbine. Darker bins indicate more frequent wind directions, lighter ones less frequent ones, thereby visualizing the wind turbine interactions. Next to the domain orientations, Figures 4.9 (a-c) also depict information relevant to all three TFs; (a) the turbine indexes, (b) the simulated domain size, and (c) the normalized distance between turbine T0 and the other turbines.

To link the dynamics back to the layout, time is also given in convective timescales. This denotes the time taken by a particle to travel a characteristic length within the domain. We choose this length to be five turbine diameters, as this is closely related to the spacing of the turbines; see Figure 4.9(c). The




Figure 4.9.: Collection of the three simulated LES TFs of the 10-turbine subset of the HKN wind farm. Panels (a-c) feature pizza-shaped histograms of the wind direction centered in the turbine locations: darker colors indicate more frequent wind directions and, therefore, turbine interactions that happen more frequently during the TF. Additionally, (a) depicts the turbine indexes; (b) the simulated domain size; and (c) the relative distance between turbine T0 and the other turbines, normalized by turbine diameters. The domains are rotated such that the initial wind direction is aligned with the precursor, and the remaining wind direction time series can be simulated with the same inflow planes.

4

freestream velocity is used to normalize the characteristic length:

$$t_c = \frac{5 \cdot 178.3 \text{ m}}{8 \text{ ms}^{-1}} = 111.4 \text{ s}. \quad (4.11)$$

4.4.2. POWER GENERATED

The power generated by SOWFA is calculated based on an actuator disc model. Simulated on a coarse grid, these tend to overestimate the power generated by the turbines, which is a known issue [100, 101]. The resulting mean ratio between the power generated in SOWFA and OFF is 1.34. Based on this mismatch, the power measurements by SOWFA in the following plots are either normalized or marked with a “c”, which denotes that the power was divided by the correction factor. Next to the LES data, the zero-phase-filtered power output data from the LES are also used to analyze model and controller performance. This filtering removes the influence of turbulence on turbine power, isolating the underlying trends more consistently with the wake dynamics that OFF aims to describe. To this end, a fourth-order Butterworth filter is used with a cutoff frequency of 1/370 Hz. The cutoff frequency is motivated by the results presented later in Figure 4.13 (b). Note that the individual turbine signals are filtered. Derivatives, like farm power or energy, then use either the filtered turbine power or the original signal and are marked with “lpf” if they use the filtered data.

The match between OFF and SOWFA is investigated in three ways: (i) on a selected turbine level for a selected controller, (ii) on a farm level for a selected controller, and (iii) on a statistical level. Figure 4.10 and 4.11 investigate the data collected for turbine T3. The data were recorded using the dead-band LuT and baseline controllers with $\phi_{lim} = 5 \text{ deg}$, and $k_i = 0.02$, one of the settings selected for validation based on the results in Figure 4.7. Turbine T3 is selected as it acts as an upstream turbine in TF 1, see Figure 4.9, and as a downstream turbine in TFs 2 and 3. This is mirrored in Figure 4.10, where the turbine produces its maximum power during the initial hours of the time series. The LuT-controlled case diverges




Figure 4.10.: Power generated (a) and efficiency with respect to the baseline (b) of turbine T3 throughout the full simulated wind direction time series. "LES c" refers to the corrected SOWFA data, and "lpf" refers to the zero-phase low-pass-filtered data. The controller settings are $\phi_{lim} = 5$ deg, and $k_i = 0.02$. The detailed data from the time frames are provided in Figure 4.11.

as the turbine engages in yaw steering and sacrifices power to redirect its wake. During later periods of the simulation, T3 becomes a downstream turbine and its power generated significantly decreases. Here, we can see an inverse effect, where T3 benefits from the yaw steering of other turbines and generates more power in the controlled case than in the baseline case.

Zooming in on the TFs that underwent LES, Figure 4.11 gives a more detailed look into the match of the LES data and the OFF data. Qualitatively we observe an overall fitting trend between the LES signal and the power predicted by OFF. An immediate difference between the two is the influence of turbulence on the LES signal. This causes noticeable variations that OFF cannot predict. The low-pass-filtered signal removes this discrepancy partially and shows a signal that is overall better aligned with the OFF signal. One aspect that gets lost due to this filtering is the response of the turbine power to yaw angle changes: Figure 4.11(b) shows the efficiency of the turbine during a period where turbine T3 engages in yaw steering to lessen the wake interaction with a downstream turbine. In OFF, the rotor misalignment causes sharp decreases and increases in efficiency, while the change is either smoothed out by filtering or hidden in the noise for the LES data. Reoccurring discrepancies between OFF and the low-pass-filtered LES data appear in the form of a phase shift, mainly visible with the baseline power signal: OFF displays slightly delayed reductions and recoveries compared to SOWFA. This might be the product of a wake that advected too slowly, which is notable as similar models specifically slowed their advection speed down for a better match with reference data. Another difference between OFF and the LES data is visible in the turbine efficiency displayed in Figure 4.11(d) and (f): OFF tends to either match or overestimate the effect of yaw steering on the turbine efficiency, compared to the filtered LES signal. This may be attributed to the fact that OFF describes a middle ground between an overconfident steady-state model and a more realistic LES.

Figure 4.12 moves from the turbine power described previously to the farm level. As the scale increases, the differences between the signals decrease. On




Figure 4.11.: Power generated (a,c,e) and efficiency with respect to the baseline (b,d,f) of turbine T3 during the three simulated TFs. "LES c" refers to the corrected SOWFA data, and "lpf" refers to the zero-phase low-pass-filtered data. The data are a subset of Figure 4.10.

a farm level OFF shows a qualitatively better match than on a turbine scale, where differences become much more clear. The farm power efficiency is also more balanced compared to the turbine level; both over- and underestimations are present if there is a mismatch, which suggests a lower bias. The improved performance on a farm scale may stem from different sources: (i) The fact that turbines are distributed throughout the farm makes it more likely that if one is not waked, another one may be. As a result under- and overestimation may cancel out. (ii) Looking at an individual turbine, small increases in wind speed lead to a large amplification of the power generated. As a result, mismatches create a large error. However, in the presented farm context the power contribution of waked turbines is small compared to the free-stream turbines. The data presented in Figure 4.12(a) and (b) also highlight TF 1 as a difficult period for wake steering to achieve consistent gains.




Figure 4.12.: Farm power generated (a,c,d) and efficiency with respect to the baseline (b,d,f) during the three TFs. "LES c" refers to the corrected SOWFA data, and "lpf" refers to the zero-phase low-pass-filtered data. The dead-band controller settings are $\phi_{lim} = 5$ deg and $k_i = 0.02$.

4.4.3. POWER SIGNAL CORRELATION

The results presented in Figure 4.10 - 4.12 show the similarities but also the discrepancies between OFF and the LES with respect to power generated. In the OFF environment, fluctuations in the power signal are due to (i) wind direction changes, (ii) control set point changes, and (iii) delayed wake dynamics. By contrast, the LES environment also reflects fluctuations due to turbulence and wake meandering. These latter two factors contribute to higher-frequency effects, raising the question: which frequency ranges does OFF effectively capture? And which frequencies could also be represented in a steady-state model?

To answer this question, we investigate the correlation between the power signals. Assuming that the discrepancies between OFF and the LES are of a high-frequency nature, one would expect that the correlation between the two models increases as high-frequency fluctuations are filtered out. In turn, with too aggressive filtering, the correlation should eventually decrease as the LES signals

lose components described by OFF. Based on these assumptions, the individual turbine data of TFs 1-3 for the baseline and LuT dead-band ($\phi_{\text{lim}} = 5$ deg, $k_i = 0.02$) controllers are correlated between OFF and the LES. A total of 180 h of data, or 18 h per turbine, are subsequently processed. Figure 4.13 (a) illustrates the influence of the cutoff frequency of the fourth-order Butterworth filter applied to the LES on the correlation score recorded by OFF while Figure 4.13 (b) depicts the resulting average correlation error. The average correlation error is defined as the mean distance of the turbines to 1 for all three TFs:

$$e_{\text{corr}} = \frac{1}{n_{\text{DT}}} \sum_{i_{\text{TF}}} \sum_{i_T} [1 - \text{corr}(p_{\text{OFF}}, p_{\text{LES}})], \quad (4.12)$$

where p is the power of turbine i_T in TF i_{TF} , and $n_{\text{DT}} = 5 + 6 + 6 = 17$ is the total number of downstream turbines considered summed across all three TFs. Combining the baseline and controlled cases, the minimum for e_{corr} is achieved for $f_{\text{cutoff}} = 1/370$ Hz = $1/3.33 t_c = 0.0027$ Hz. In contrast to OFF, the collective minimum for FLORIS is reached at $1/520$ Hz = $1/4.68 t_c^{-1} = 0.0019$ Hz, so at a lower frequency. This gap is explained by the added wake dynamics in OFF, as OFF uses the same FLORIS model in its core. This cutoff also aligns with the literature on wake meandering, which is not captured by OFF: [102] finds the wake meandering frequency to be around $\frac{u_{\infty}}{20D}$, which equals 0.0022 Hz for the presented study. [17], on the other hand, suggest a higher frequency, which, for this study, equals 0.022 Hz. We can conclude that OFF does describe the wake dynamics up to the wake meandering frequency. Additionally, we note that OFF leads to a lower error than FLORIS; while OFF finds its minimum at $e_{\text{corr}} = 0.11$, FLORIS returns $e_{\text{corr}} = 0.19$. It should be noted that the filtering timescale used to preprocess the wind direction signal (1/600 Hz) may limit OFF's performance, as it filters out relevant dynamic scales. Related work by [103] suggests, for instance, that mean wind direction changes may occur with a frequency of up to 1/270 Hz. Rerunning the LES with a higher cutoff frequency would likely increase OFF's effective cutoff frequency estimation; however, this was not feasible within the scope of the present work.

Figure 4.14 provides more insight into the source of the correlation error. Figure 4.14 (a) and (b) show the correlation error in OFF, split into LuT cases (a) and BL cases (b). This is accompanied by the results for FLORIS, depicted in (c) and (d), also split into LuT cases and BL cases, respectively. Upstream turbines, like T0, T1, T3, T5, and T7 for TF 1, are neglected in Figure 4.13 and 4.14 as they are operating at close-to-maximum power in OFF and FLORIS, while their LES counterparts are affected by turbulence, see for instance Figure 4.11 (a). As a result, the turbines modeled in OFF and FLORIS experience no excitation, while the LES ones do. This leads to effectively no correlation between the signals.

Looking at which turbines lead to the larger e_{corr} for FLORIS, the turbines in TF 1 contribute a large share, as well as turbine T9 in TF 2. Based on Figure 4.9, we can see that TF 1 features long-distance turbine-to-turbine interactions. This fact, paired with the varying wind direction, leads to a situation where the steady-state approximation of FLORIS fails and where wake dynamics play a significant role in the power generated. This also complements the observation from Figure 4.12 (b), where it was visible that TF 1 is a challenging case for the steady-state-based LuT controller. A notable similarity between OFF and FLORIS is that the LuT cases lead to a higher error than the baseline cases. One reason for this discrepancy




Figure 4.13.: (a) Correlation of the downstream turbine power in OFF and the LES. The LES data are zero-phase low-pass-filtered with varying cutoff frequencies. Each line corresponds to the correlation of one turbine, blue lines are the data from the controlled TFs, and red lines are from the baseline ones. The dot represents the maximum correlation from a given turbine. The average error is depicted in (b) and is minimal for $f_{\text{cutoff}} = 1/370 = 0.0027$ Hz. Additionally, there is the line for the correlation of the FLORIS data with the LES. Its minimum is located at $1/520$ Hz. The dead-band controller settings are $\varphi_{\text{lim}} = 5$ deg and $k_i = 0.02$.

could be that the turbine model does not accurately capture the impact of larger misalignment angles. This would motivate turbine model corrections as suggested by [104] and [89]. Additionally, this error may be partially rooted in the wake dynamics triggered by LuT control. Indeed, LuT-based wake steering tends to amplify changes in wind direction: a variation of just a few degrees in the wind direction may, under certain circumstances, induce a yaw-offset angle change that is 10 times greater than the original wind direction change [99]. This results in more frequent and larger variations in wake states.

4.4.4. POWER ERROR STATISTICS

Section 4.4.2 first investigates the turbine power, then the farm power, as well as the role of timescales. This discussion is limited to one set of controller settings $\varphi_{\text{lim}} = 5$ deg and $k_i = 0.02$. For brevity, we denote the controller settings as $(\varphi_{\text{lim}}, k_i)$ in the following paragraph. Two more sets of settings underwent LES, namely $(2, 0.09)$ and $(10, 0.05)$. Table 4.1 summarizes characteristic error quanti-




Figure 4.14.: Cumulative correlation error between the turbine power from the LES and OFF (a,b) and FLORIS (c,d). The data are split into the LuT cases (a,c) and the baseline cases (b,d). The shaded areas indicate the contribution of each downstream turbine across the three TFs on top of each other. With (d) it is indicated which layer relates to the corresponding TF. The dead-band controller settings are $\phi_{lim} = 5$ deg and $k_i = 0.02$.

ties for all controllers. The table combines the three TFs for each controller setting to calculate the difference between the OFF prediction and FLORIS prediction. The table lists the normalized RMSE for the turbine and farm power, as well as the correlation of both signals. The normalization was done with the corrected LES data. The values show that the addition of dynamics renders OFF more robust towards the addition of yaw steering, compared to FLORIS: while the turbine and farm NRMSE slightly decrease for OFF, there is a notable increase for FLORIS related to the switch from BL to LuT operation. Similarly, the correlation of the farm and turbine power decreases for both OFF and FLORIS, but the steady-state approximation results in a larger decrease; e.g., for (2, 0.09), the farm power correlation by OFF decreases by ≈ -0.03 compared to ≈ -0.11 for FLORIS. However, both OFF and FLORIS achieve similar correlation and error results for baseline operation. An explanation can be that the LuT creates wind farm states that are more sensitive to environmental changes. As a result, the modeled wake dynamics become more relevant. Also notable is the NRMSE decrease for both models with the switch from turbine level to farm level, from values between 0.17 and 0.27 to values between 0.04 and 0.06. Consequently, model inaccuracies on a turbine level do not necessarily lead to equally large errors on a farm level. This also indicates that going forward, improved model descriptions might lead to less uncertainty on a turbine basis but might show diminishing returns on a farm scale.

Model	ϕ_{lim}	k_i	Mode	T. NRMSE [-]	T. Corr. [-]	F. NRMSE [-]
OFF	2	0.09	LuT	0.19	0.81	0.047
FLORIS	2	0.09	LuT	0.27	0.74	0.064
OFF	2	0.09	BL	0.20	0.88	0.048
FLORIS	2	0.09	BL	0.19	0.87	0.045
OFF	5	0.02	LuT	0.18	0.83	0.043
FLORIS	5	0.02	LuT	0.24	0.76	0.056
OFF	5	0.02	BL	0.20	0.88	0.047
FLORIS	5	0.02	BL	0.20	0.87	0.045
OFF	10	0.05	LuT	0.18	0.85	0.042
FLORIS	10	0.05	LuT	0.24	0.80	0.053
OFF	10	0.05	BL	0.20	0.88	0.048
FLORIS	10	0.05	BL	0.21	0.86	0.047

Model	ϕ_{lim}	k_i	Mode	F. Corr. [-]	f_{cutoff} [Hz]	e_{corr} [-]
OFF	2	0.09	LuT	0.88	1/360	0.14
FLORIS	2	0.09	LuT	0.81	1/540	0.26
OFF	2	0.09	BL	0.90	1/430	0.10
FLORIS	2	0.09	BL	0.92	1/520	0.13
OFF	5	0.02	LuT	0.90	1/360	0.12
FLORIS	5	0.02	LuT	0.84	1/520	0.24
OFF	5	0.02	BL	0.91	1/390	0.10
FLORIS	5	0.02	BL	0.92	1/520	0.15
OFF	10	0.05	LuT	0.91	1/370	0.11
FLORIS	10	0.05	LuT	0.85	1/510	0.21
OFF	10	0.05	BL	0.91	1/370	0.09
FLORIS	10	0.05	BL	0.91	1/500	0.16

Table 4.1.: Power error statistics for each controller tested in OFF, FLORIS and LES. From left: T.NRMSE, the normalized root-mean-square error calculated with the corrected turbine power LES data; T.Corr., the correlation with the unfiltered turbine power LES signal; F.NRMSE, the normalized root-mean-square error calculated with the corrected farm power LES data; F.Corr., the correlation with the unfiltered farm power LES signal; f_{cutoff} , the cutoff frequency for LES filtering; and e_{corr} , the average correlation error.

4.4.5. ENERGY GENERATED

Section 4.4.2 investigates the power generated by the wind farm at different time and turbine scales. This section complements the results with a discussion about the energy generated. More specifically, the efficiency of the wind farm is compared between the LES and the surrogate models. The efficiency is calculated as the ratio of the farm energy generated using LuT control, normalized by BL control, integrated over a time window ΔT :

$$\eta(t, \Delta T) = \frac{\int_t^{t+\Delta T} \Delta t \sum_{n_T} p_{LuT}(\tau) d\tau}{\int_t^{t+\Delta T} \Delta t \sum_{n_T} p_{BL}(\tau) d\tau}, \quad (4.13)$$

where p refers to the power generated by a turbine, Δt is the time step, and t is the time. Figure 4.15 compares $\eta_{LES}(t, \Delta T)$, the wind farm efficiency simulated




Figure 4.15.: Wind farm efficiency as predicted by the surrogate models OFF (a-d) and FLORIS (e-h) and as simulated in the LES. The efficiency is calculated based on the ratio of energy generated over a time window ΔT , which is equal for each column of the figure, e.g. (a) and (e). The dotted white line indicates a perfect fit, which is complemented by the linear regression of the data, given as red line and equation. The color map is normalized by the largest bin count based on the given time window. The darkest color is reserved for the smallest non-zero bin count; empty bins are not filled. Note that the distribution of ΔT is not equidistant.

in the LES, with $\eta_{\text{OFF}}(t, \Delta T)$ and $\eta_{\text{FLORIS}}(t, \Delta T)$, the values OFF and FLORIS predict respectively. This is done for four values of ΔT between 100 s and 1800 s with data from all three TFs, and based on the $\phi_{\text{lim}} = 5$ deg, $k_i = 0.02$ controllers. A first observation is that the range of values for the farm efficiency decreases with increasing length of ΔT . This shows the increasing convergence towards a more consistent controller performance over a longer time as well as a diminishing influence of effects at a small timescale. In comparison, between OFF and FLORIS, OFF generally predicts a narrower fit for small values of ΔT , closer to the ideal correlation line. With increasing ΔT , this difference diminishes, and the distributions of FLORIS and OFF become more equal. For large ΔT , FLORIS shows a structural underestimation compared to the LES data, where OFF still predicts values along the ideal correlation line. This observation is also quantifiable with the linear regression parameters: as ΔT lengthens, the linear coefficient approaches 1, and the bias decreases. This trend is visible for both models; however, OFF consistently presents parameters closer to the ideal values.

Figure 4.16 investigates the error in the approximation of the farm efficiency to further quantify and compare the differences. For each TF and each simulation environment $\eta(t, \Delta T)$ is calculated for $\Delta T \in [100, 1900]$ s and $t \in [t_0, t_1 - T]$, where t_0 is the start time of each TF and t_1 is the final time. Figure 4.16 compares how the root-mean-square error between the $\eta(t, T)$ from the LES and the $\eta(t, \Delta T)$




Figure 4.16.: Root-mean-square error in the wind farm efficiency in LES compared to OFF or FLORIS. The wind farm efficiency is defined as the ratio of the energy generated with LuT control divided by the baseline energy integrated over a given time window.

4

of OFF and FLORIS changes for different T . The difference between the LES and FLORIS improves significantly for longer averaging periods, highlighting its design meant for long-term wind farm behavior. On the other hand, OFF benefits from the addition of wake dynamics and shows much lower RMSE values compared to FLORIS. However, this advantage becomes smaller as ΔT grows larger. As a result, a user has to decide if the added computational cost of OFF in comparison to FLORIS justifies the improvement in prediction. We conclude that, based on this case study, it is advantageous to use OFF for quantities of interest shorter than ≈ 20 min. However, for longer timescales the benefit of the added dynamics diminishes.

4.4.6. COMPUTATIONAL COST

One of the main motivations for dynamic parametric wake models like OFF, or by extension for FLORIDyn or OnWaRDS, is the low computational cost compared to high-fidelity numerical methods such as LES, for instance. On the other hand, it is evident that the computational cost has to be higher than the cost of the underlying steady-state wake model. Simplified, the computational cost of both OFF and time-marching FLORIS can be expressed as a function of the number of time steps n_k , the number of turbines n_T , and the number of observation points n_{OP} :

$$\begin{aligned} \mathcal{O}_{\text{OFF}} = n_k \cdot \dots & \\ \left[\underbrace{\mathcal{O}_{\text{State prop.}}(n_T, n_{OP}) + \mathcal{O}_{\text{TWF}}(n_T, n_{OP}) + n_T \cdot [\mathcal{O}_{\text{F. run}}(n_T) + \mathcal{O}_{\text{F. reinit.}}(n_T)]}_{\mathcal{O}_{\text{prediction}}} + \dots \right. & \\ \left. \dots \mathcal{O}_{\text{corr.}}(n_T, n_{OP}) + \mathcal{O}_{\text{con.}}(n_T) \right] + \mathcal{O}_{\text{F. init.}} + \mathcal{O}_{\text{OFF init.}}, & \quad (4.14) \end{aligned}$$

$$\mathcal{O}_{\text{FLORIS}} = n_k \cdot \mathcal{O}_{\text{F. run}}(n_T) + \mathcal{O}_{\text{F. init.}} + \mathcal{O}_{\text{con.}}(n_T), \quad (4.15)$$

where $\mathcal{O}_{\text{State prop.}}$ refers to the cost of the state propagation, \mathcal{O}_{TWF} to the creation of the TWFS, $\mathcal{O}_{\text{F. run}}$ to the cost of one FLORIS evaluation, $\mathcal{O}_{\text{F. reinit.}}$ and $\mathcal{O}_{\text{F. init.}}$ to the FLORIS reinitialization and initialization, $\mathcal{O}_{\text{corr.}}$ to the state correction, and lastly $\mathcal{O}_{\text{con.}}$ to the derivation of the control set points. This is accompanied by other costs, such as visualization, data storage and memory limitations, which are excluded here.

Performance analysis during the code development has shown that the re-occurring computational costs of $\mathcal{O}_{F, \text{reinit.}}$ can be substantial depending on the implementation. FLORIS was developed with other simulation goals in mind. This leads to costs associated with the reinitialization that are mandatory for some FLORIS applications but could be neglected for purposes of the OFF simulations. Consequently, existing codes similar to OFF have mainly chosen to implement their own wake model. This, in return, limits the capabilities and flexibility of the wake model, which was one of the main motivations for the development of OFF. Another consideration to reduce computational costs is to only run relevant turbines in the steady-state simulation and thereby decrease the cost of $\mathcal{O}_{F, \text{run}}(n_T)$. This could be done by excluding turbines that do not contribute to the wake losses experienced by the turbine the TWF is dedicated to. The validity of this approach also depends on the steady-state model capabilities. For instance, if there is a blockage model based on n_T , this simplification would introduce a systematic model error. Lastly, parallelization is a natural approach to improving computational complexity. The n_T TWF evaluations done in one time step can be done independently of one another, which would lead to a performance improvement for up to n_T cores. In this work, we investigate a large number of control settings and, therefore, use OFF as a single-core code and split the task at hand over multiple simulations. To give an estimate, in our 10-turbine simulations, OFF ran with a real-time factor of $2.2 \cdot 10^{-1}$ in single-core performance, resulting in 5 h 20 min CPU time for 23 h 45 min simulated time. The SOWFA simulations, recalculated from 80 cores to 1 core, ran with a real-time factor of $2 \cdot 10^3$, resulting in 6030 h CPU time for 3 h simulated time. Lastly, FLORIS ran with a real-time factor of $5.2 \cdot 10^{-5}$, resulting in 4.43 s wall time for 23 h 45 min simulated time. Previous work showed that the real-time factor of a model like OFF can be reduced to the order of 10^{-3} for a similar-sized wind farm with a dedicated implementation of the Gaussian wake model [47, 75].

4.5. CONCLUSION

This paper introduces OFF, a dynamic open-source wake model designed for wind farm flow control and wake model development, and as a unified interface for various similar models. In this context, a generic description of a passive Lagrangian particle wake model is provided, along with details on the specific version used to achieve the results discussed here. In an example case, the model is used to make an informed parameter choice for a wake steering controller before verifying the selected settings in LES. The controller applies a wake steering look-up table dynamically for a 10-turbine wind farm. The wind farm layout is based on the Hollandse-Kust-Noord wind farm, and the approximately 24 hour long period of wind direction time series used to test the controllers is based on field data from the same location.

The results from the study show that the wind farm controller can lead to suboptimal performance in the presence of wind direction changes compared to what was predicted during the generation of the LuT based on a steady-state assumption. The study also shows that the wake steering controller's performance can vary widely for the same wind direction based on the prior state of wind direction, wakes and controllers used. Six selected sets of controller settings are then verified in LES in three 3 hour long subsets of the wind direction change time series. The results show overall good agreement between the LES and OFF in

both predicted power generated and wake steering controller efficiency. The LES, for instance, confirms that one of the selected time frames creates a challenging environment for the wake steering controller to return consistent gains over the baseline operation. The results further investigate the timescales described by both FLORIS and OFF. A conclusion drawn from the comparison is that a dynamic wake description leads to a better correlation with the LES power signal, as well as a lower root-mean-square error compared to a steady-state prediction.

In conclusion, OFF provides a unified interface to a dynamic wake description that is advantageous over steady-state wake models for shorter time periods (< 20 min). The model is open-source and designed to interface with steady-state wake model toolboxes. This has been demonstrated with the FLORIS toolbox. As a result, users of OFF can also benefit from the ongoing development done for the underlying wake models.

4

Future work should further investigate the use and effect of various steady-state wake models in a dynamic context. This starts with further validation of the approach and the generation of more realistic test and reference cases. One short-coming of the presented case study is its limitation to wind direction variations. Future work should investigate the model and control performance with realistic wind speed variations, similar to the works of, for example, [51, 95]. It may also involve investigating the selection of wake parameters. Since OFF describes wakes at higher frequencies, the resulting wake shape may appear more slender than a steady-state wake, which must account for small-scale wind direction changes and wake meandering. The OFF code is further built in a modular way to be expanded by other dynamic elements and to further explore their effectiveness for the description of dynamic flows. This includes, for instance, wake advection descriptions (e.g. [75, 85, 105]), shear and veer parameterizations (e.g. [106]), or floating turbine dynamics (e.g. [15]). Another direction of interest can be the employment of single-wake dynamic surrogate models in a wind farm, e.g. [107, 108].

In the long term OFF should lead towards a new dynamic wake model that replaces modularity with reduced computational cost and a dedicated, informed selection of the components previously explored.



5

ENSEMBLE BASED flow field ESTIMATION USING THE DYNAMIC WIND FARM MODEL FLORIDYN

This chapter is based on the following publication:

[83] M. Becker, D. Allaerts, and J.W. van Wingerden, *Ensemble-Based Flow Field Estimation Using the Dynamic Wind Farm Model FLORIDyn*, *Energies* (2022), 10.3390/en15228589.

Wind farm control methods allow a more flexible use of wind power plants over the baseline operation. They can be used to increase the power generated, to track a reference power signal or to reduce structural loads on a farm wide level. Model based control strategies have the advantage that prior knowledge can be included, for instance by simulating the current flow field state into the near future to take adequate control actions. This state needs to describe the real system as accurately as possible. This paper discusses what state estimation methods are suitable for wind farm flow field estimation and how they can be applied to the dynamic engineering model FLORIDyn. In particular, we derive an Ensemble Kalman Filter framework which can identify heterogeneous and changing wind speeds and wind directions across a wind farm. It does so based on the power generated by the turbines and wind direction measurements at the turbine locations. Next to the states, this framework quantifies uncertainty for the resulting state estimates. We also highlight challenges that arise when ensemble methods are applied to particle-based flow field simulations. The development of a flow field estimation framework for dynamic low-fidelity wind farm models is an essential step toward real-time dynamic model based closed-loop wind farm control.

5.1. INTRODUCTION

Wind turbines are frequently positioned in a wind farm in such a way as to minimize electrical cabling costs and to utilize given space as efficiently as possible. In this context, wake effects can have a significant impact on the power generated by the wind farm: as one turbine extracts energy from the flow, a downstream turbine experiences lower wind speeds and generates less energy [4]. The impact of wakes can be reduced by wind farm layout optimization and by changing the way wind turbines are controlled. Wind farm control (WFC) methods utilize the degrees of freedom a wind turbine provides in order to improve the wind farm's performance [5]. This can be done for instance to increase power generated or to reduce loads on the turbines to extend their lifetime.

Model based WFC methods utilize a surrogate model to find the optimal control inputs for the wind farm. This has shown to be successful in steady state, for instance with the method of wake steering [4], where turbines are purposefully misaligned with the main wind direction to deflect their wake. For steady state conditions, set points can be calculated offline and are then applied using a look-up table during operation. More recent publications iterate on the approach and aim to also include dynamic effects, some of which are discussed here:

Dynamic wake effects can be approximated at a low computational cost by reducing the underlying physics. One example is the free-vortex method, which reduces the simulation scope to vortices shed by the rotor and their interaction with one another. Recent work has shown that the free-vortex method, paired with an adjoint-optimisation, can be used for model predictive control [109]. In the presented work, the algorithm is applied to a two-turbine wind farm during a wind direction change, where the flow field conditions are prescribed on a global scale. While successful at a small scale, the method is currently limited to low numbers of turbines. This is due to a steep increase in computational cost and numerical instabilities. The free-vortex method further has inherent difficulties to accurately describe the wake behavior once the wake does break down. The distance at which the free vortex method begins to become inaccurate is subject to ongoing research and can vary across different implementations. Another way of approximating dynamic effects in wind farms is to decrease the time scale at which steady state models are used. In [110] the authors use a steady state model but adapt it to take changing atmospheric conditions on a longer time scale into account (minutes to hours). This way they achieved dynamic closed loop control with a steady state model across a full diurnal cycle. The model does not feature a dedicated flow field model, but rather averages the environmental conditions over time. On a shorter time scale of seconds to minutes, there are only a few models which aim to simulate the flow behavior in a wind farm at a low computational cost in order to provide control inputs in real-time.

One such model is the FLOW Redirection and Induction Dynamics (FLORIDyn) model, originally published in [19] and more recently revised in [1] and [47]. The model can take heterogeneous and transient flow conditions into account and can simulate the propagation of turbine state changes through the wake. This is achieved by creating particles, so called Observation Points (OPs), at the rotor plane. These inherit the turbine state and the wind field state at the time of their creation. Each OP then proceeds to propagate downstream according to the wind speed and direction it has inherited. The turbine state, together with the wind field state and the OP location, provide enough information to

approximate the influence of the turbine wake in the proximity of the OP. The FLORIDyn model uses a parametric wake model, designed for steady state, to calculate the wind speed deficit. The model therefore adds flow and advection dynamics to a computationally cheap steady state wake model. Multiple similar models have been published recently, which are briefly described below.

The *FOWFSim-Dyn* model presented by [15] provides a basis for a dynamic description of floating wind farms. It couples a parametric wake model with turbines on floaters, constrained by mooring lines. Based on the inflow, the turbines change their location and the wakes adapt. The inflow is assumed to be uniform throughout the domain and is modeled as an imposed function of time. The model *UFLORIS* is presented in [76]. It makes similar design choices as FLORIDyn, in its switch from [1] to [47]: instead of using multiple OP chains to cover the entire wake, the model employs only one chain of OPs along the center line. *UFLORIS* employs a 2D wake and models the wind speed as part of the Observation Point's state. The model presented in [75] takes a different approach and models the wind field as its own set of ambient OPs which propagate at a different speed than the wake OPs. The model also incorporates crosswind components at the rotor plane and is able to show meandering effects in the wake. Similarly, [74] also employs a set of wake-OPs and ambient-OPs to differentiate between the background flow behavior and the wake dynamics.

The emergence of various dynamic parametric wake models shows that the field is maturing and different design choices are being explored. A common goal for these models is to be applied in a wind farm context for real-time control purposes. To achieve this, an estimation framework is needed to identify the current wind field state. The identified state can then be used to simulate into the near future and to decide on the best control actions. Without estimation, the model does not have knowledge of the real life circumstances and does not lead to useful predictions. In previous work [21], a Kalman Filter is proposed to estimate the wind speed in FLORIDyn with promising results. However, the work is lacking a wind direction estimation and this approach does not estimate the uncertainty of the system states but rather assumes prior knowledge of it.

An alternative estimation method is the Ensemble Kalman Filter (EnKF) [111]. The EnKF follows an ensemble based data assimilation approach and aims to estimate the state by simulating many different versions of the model. Each realization of the model has its own state and forms one ensemble. The states differ across all ensembles and diverge over time, if not corrected. During the correction step, the differences between the ensemble states are used to approximate the state-error-covariance matrix, which allows the calculation of the Kalman gain matrix. The capability to approximate the state-error-covariance matrix is property of the EnKF which simplifies what is a major tuning effort in other Kalman filter implementations. Therefore, the state estimation of nonlinear systems with the EnKF is relatively easy compared to other methods, which might require more system information, such as derivatives. The reduced mathematical effort comes at the cost of more computational effort in the simulations. As a result of increasing computational capacities and the parallelizable nature of the EnKF, this trade-off has become more tolerable. The method can be used to estimate the state of complex nonlinear systems, as well as the uncertainty of an identified state, a property which is useful for solving control problems in a robust manner. Successful applications of the EnKF include flow problems [112–117], which make it particularly interesting for wind farm flow field estimation.

In a similar problem setting to FLORIDyn, an EnKF is used in [114] with the wind farm model WFSim to estimate the state of the flow. The EnKF returns promising results including the uncertainties of the states in addition to a computational setup which can easily be adopted to keep the computational cost low. However, an elementary difference between the work of [114] and FLORIDyn is that WFSim has states at fixed locations, whereas FLORIDyn propagates its states and is essentially a particle simulation. This requires an adaptation on how the EnKF can be applied compared to its textbook examples. The work of [112] pioneered the application of the EnKF in an adaptive mesh simulation, something which has been further developed in [113, 115]. A recent publication [117] adapts the problem statement by incorporating Lagrangian particles in a mesh simulation. This work presents characteristics close to FLORIDyn and shows that the EnKF can be applied for particle simulations. The mere fact that an EnKF framework can be applied to a dynamic-low-fidelity model has been shown by [116]: based on the power generated an EnKF corrects the wind speed deficit and the wake expansion.

Similar estimation techniques to the EnKF are the Unscented Kalman Filter (UKF) [118], and variational data assimilation, such as the Four-Dimensional Variational method (4D-Var) [119]. The UKF propagates selected versions of the state vector, called *sigma vectors* (or sigma points). The sigma vectors are created based on the state covariance. They are then propagated in time using the system equation and the weighted mean of the resulting states is the estimate. The error covariance between the sigma vectors and the outputs is then used to calculate the Kalman Gain matrix and to correct the estimate. This allows the UKF to do state estimation for nonlinear systems [120]. The difference to the EnKF is, that the sigma vectors are then reseeded, based on the new state covariance. The EnKF on the other hand propagates the same ensembles further in time. A downside of the UKF is the number of sigma vectors, which is typically twice the number of states. For FLORIDyn, this would result in hundreds to thousands of sigma vectors. Variational data assimilation methods optimize an initial state to fit the past outputs produced by the model over an assimilation window. The identified state can then be used to predict future model behavior. In particular, 4D-Var is used with success for meteorological and flow simulations [121], similar to the EnKF. This frequently leads to the question when to choose one over the other [122]. What makes 4D-Var unsuited for FLORIDyn is that the OPs with the identified state leave the system boundaries and are disregarded. As a result, the algorithm would put effort into estimating states which do not have an influence anymore at the current time step, nor in the future. We conclude that the EnKF remains the most practical and promising approach for the simulation circumstances of this work.

To summarize, dynamic-parametric analytical wake models, such as FLORIDyn, grow in popularity. They can approximate dynamic flow behavior within a wind farm at a low computational cost, which makes them suitable for real-time closed-loop wind farm control. In order to make meaningful control decisions, the model state needs to be equal to the real wind farm state. To align the two states, an estimation framework is needed. The review of existing methodologies suggests that the Ensemble Kalman Filter is the most suitable estimation framework but requires some adaptation.

The main contribution of this paper is an Ensemble Kalman Filter framework to jointly estimate the background wind speed and wind direction in a wind

farm, using the dynamic parametric wind farm model FLORIDyn. The presented framework is innovative as wind speed and wind direction estimation are generally treated as separate problems. The work further contributes to the recent efforts to explore how the Ensemble Kalman Filter as a method can be used to estimate the states of Lagrangian particle simulations. The estimation is based on already available turbine data such as the power generated and wind vane measurements. The results and insights of this work are also relevant for other dynamic parametric wind farm models mentioned earlier [15, 74–76], as well as for other Lagrangian particle simulations.

The remainder of the paper is structured as follows. Section 5.2 discusses FLORIDyn's properties and presents the resulting Ensemble Kalman Filter framework. Results obtained with the new framework are presented in Section 5.3. Section 5.4 draws conclusions and gives an outlook for future work.

5.2. MATERIALS AND METHODS

In order to implement the proposed Ensemble Kalman Filter approach, FLORIDyn needs to be described as a state space system. We will approach this problem by first discussing the properties of the FLORIDyn algorithm and the resulting differences to other simulation types in Section 5.2.1. Following these insights, we present different ways to formulate the Ensemble Kalman Filter framework in Section 5.2.2. We will also discuss extensions of its formulation¹.

5.2.1. PROPERTIES OF THE FLORIDYN APPROACH

FLORIDyn is a particle simulation approach to model the dynamic behavior of wind turbine wakes given environmental conditions. In practice, so called Observation Points (OPs) are created at every time step and propagate downstream with the free wind speed and along the main wind direction. The term *free wind speed* refers to the assumed background wind speed, unaffected by the wakes. The OPs inherit the state of the turbine at their time of creation. This allows them to calculate the wake of the turbine at and around their location. The detailed process is described in [47]. The wind field states used for propagation (wind speed, direction and ambient turbulence intensity) are also part of the states of an OP. The full state of an OP is given by its location $\mathbf{x}_{L,OP}$, the turbine states attached to it $\mathbf{x}_{T,OP}$ and the wind field states $\mathbf{x}_{WF,OP}$. Combining all OP states, the system is then propagated as follows:

$$\begin{bmatrix} \mathbf{x}_L(k+1) \\ \mathbf{x}_T(k+1) \\ \mathbf{x}_{WF}(k+1) \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{L,L} & \mathbf{0} & \mathbf{A}_{L,WF}(\mathbf{x}_{WF}(k)) \\ \mathbf{0} & \mathbf{A}_{T,T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{WF,WF} \end{bmatrix} \begin{bmatrix} \mathbf{x}_L(k) \\ \mathbf{x}_T(k) \\ \mathbf{x}_{WF}(k) \end{bmatrix} + \dots \\ \begin{bmatrix} \delta(\mathbf{x}_L(k), \mathbf{x}_T(k), \mathbf{x}_{WF}(k)) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{B}_L & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_{WF} \end{bmatrix} \begin{bmatrix} \mathbf{I}_T(k) \\ \mathbf{x}_{T,0}(k) \\ \mathbf{x}_{WF,0}(k) \end{bmatrix}, \quad (5.1)$$

¹The mathematical notation of the paper is as follows: italic, non-bold letters denote scalars (e.g. x, C_p), bold, lowercase letters denote column vectors (e.g. $\mathbf{x}_{OP}, \boldsymbol{\mu}_{WF}$), column vectors with hat denote state estimates (e.g. $\hat{\mathbf{x}}_{WF}$), with bars averages (e.g. $\bar{\mathbf{x}}_{WF}$). Bold, uppercase letters denote matrices (e.g. \mathbf{A}, \mathbf{C}), matrices with tilde represent matrices modified by weighting or localization (e.g. $\tilde{\mathbf{A}}, \tilde{\mathbf{C}}$). Square brackets organize equations or define matrices and vectors, round brackets are function inputs, properties or units.

where \mathbf{x}_L refers to all OP location states, \mathbf{x}_T and \mathbf{x}_{WF} to all stored turbine states and the stored wind field states respectively. The matrices $\mathbf{A}_{L,L}$, $\mathbf{A}_{T,T}$ and $\mathbf{A}_{WF,WF}$ have a similar lower-diagonal block structure which ensures that one state is propagated to the next row for each turbine and its OPs. The propagation of the OPs following the main wind direction and wind speed is described by the matrix $\mathbf{A}_{L,WF}(\mathbf{x}_{WF}(k))$. The nonlinear term $\delta(\mathbf{x}_L(k), \mathbf{x}_T(k), \mathbf{x}_{WF}(k))$ describes the centerline deflection, as the model presented in [47] describes only OPs moving along the centerline. This term is zero in the presented work, as yaw-misalignment is not part of the later presented case studies. A detailed description can be found in [16]. As inputs there are the turbine rotor locations $\mathbf{l}_T(k)$, the current turbine state $\mathbf{x}_{T,0}(k)$ and the current wind field state at the turbine locations $\mathbf{x}_{WF,0}(k)$. The turbine rotor locations are used to determine where the OPs are created, the remaining inputs serve as information stored in those newly created OPs. The full system is described in more detail in Appendix 5.A.1.

In order to estimate the wind field states, they have to be observable. This depends on the way the output of the simulation is defined. We assume that we can only measure at the locations of the turbines and treat the turbines as sensors in our framework. For the wind direction, we assume to have a wind vane available at the turbine location. To estimate the background wind speed, we can utilize the power generated. This way the algorithm requires the least input from the turbine and utilizes already available data. On the FLORIDyn side, the power generated is estimated by calculating the effective wind speed at the turbine location and the power coefficient C_P :

$$P_{\text{turbine}} \stackrel{!}{=} \hat{P}_{\text{FLORIDyn}} = \frac{1}{2} A_{\text{rot}} \rho C_P \hat{u}_{\text{eff}}^3 \quad (5.2)$$

$$\hat{u}_{\text{eff}} = \hat{u}_{\text{free}} \prod_{i=1}^{n_T} [1 - r_i], \quad (5.3)$$

where \hat{u}_{free} is the estimate of the free / background wind speed, stored in the OPs, and all other variables with hat symbol are derived estimates. The rotor area is A_{rot} , ρ denotes the air density and r_i the wind speed reduction by the i -th wind turbine wake. The power coefficient $C_P = 4a(1-a)^2$ is calculated based on the actuator disc theory, in this work solely based on the axial induction factor as yaw is not part of the study [123]. We assume greedy control for all turbines and set a to the Betz limit of $1/3$. The wind speed reduction is based on the Gaussian wake model [16] and implemented as described in [1]. Note, that the power generated is dependent on \hat{u}_{eff}^3 , which makes an estimation of the wind speed, and its uncertainties, even more relevant. Equation (5.3) uses the free wind speed at the turbine locations, which is based on the OP state \mathbf{x}_{WF} . It is necessary to understand how all states are connected to the few estimates at the turbine locations to ensure observability. In the model presented in [47] the states of the closest upstream and downstream OP would be used to interpolate the free wind speed. This was motivated by a similar approach in the initial FLORIDyn implementation [19]. If we define a vector $\hat{\mathbf{u}}_{\text{free}}$ as (intermediate) output of our system at every turbine location, the output matrix could look as follows:

$$\hat{\mathbf{u}}_{\text{free}} = \mathbf{C}(k) \mathbf{x}_{\text{WF},u} \quad (5.4)$$

$$\mathbf{C}(k) = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & w_1 & w_2 & 0 & \cdots & 0 \end{bmatrix}, \quad (5.5)$$

where $\mathbf{C}(k)$ is the time varying output matrix of the state space system and w_1, w_2 are the non-zero interpolation weights. Note that Equation (5.5) is formulated for a two turbine case where the first turbine is uninfluenced and the second one is in the wake of the first one. The free wind speed at the second turbine is estimated from two OPs in the wake of the first turbine, based on the weights w_1 and w_2 . Note that we only use the velocity entries of \mathbf{x}_{WF} . This formulation creates a very sparse \mathbf{C} matrix. Figure 5.1 visualizes this by coloring the OPs which are used to determine $\hat{\mathbf{u}}_{\text{free}}$.




Figure 5.1.: Interaction between OPs and a turbine: the initial case only considered the two closest OPs (upper figure), while the weighted case considers a broader range of OPs (lower figure) to estimate the wind field states at the downstream turbine location.

It is here where the inherent issue lies of applying the Ensemble Kalman Filter Framework. The EnKF will be discussed in more detail in Section 5.2.2, but the main idea is that the estimator works by employing multiple versions of the model, all evolving slightly differently over time. As this is the case, the wind speed states of different ensembles diverge and so do the locations of the OPs. Since the formulation of \mathbf{C} is so sparse, different OPs and therefore different states contribute to the wind speed estimate in each ensemble. The estimator framework is built around the premise to estimate and correct the same state across all ensembles, which is no longer the case. Therefore, the formulation of \mathbf{C} has to be altered and its non-zero entries have to be wider spread, which leads to more robustness across different ensembles. Furthermore, it has to be ensured that the same states are identified across all ensembles. To solve this task a method from a very similar model is utilized: Lejeune et al. [75] employ a spatio-temporal averaging approach to weight the wind field states. We adopt this approach in FLORIDyn by weighting every OP based on its downwind and

crosswind distance to the location of interest and based on the time passed since its creation. At the location of interest, the weighted average of all OPs' states returns the free wind speed estimate. We calculate the weights using a Gaussian function [75]:

$$w(d_{dw}, d_{cw}, t_{OP}) = \exp\left(-\left[\frac{d_{dw}^2}{2\sigma_{w,dw}^2} + \frac{d_{cw}^2}{2\sigma_{w,cw}^2}\right]\right) \exp\left(-\frac{[t - t_{OP}]^2}{2\sigma_{w,t}^2}\right), \quad (5.6)$$

where d_{dw} and d_{cw} are the downwind and crosswind distance to the location of interest, t_{OP} is the time at which the OP was created and t is the current time. Figure 5.1 indicates how a broader range of OPs is now considered. This weighting also introduces three new tuning constants $\sigma_{w,dw}$, $\sigma_{w,cw}$ and $\sigma_{w,t}$, which control the downwind, crosswind, and temporal width of the Gaussian weighting function, respectively.

This way of calculating the weights is applied in two places. First, the calculation of the \mathbf{C} matrix is adapted. Its entries become larger for close and younger OPs, and become smaller for older, further away OPs. The same method is now also applied to the calculation of the propagation distance of the OPs, represented by $\mathbf{A}_{L,Wf}(\mathbf{x}_{WF}(k))$ in Equation (5.1). This change overcomes the issue of one OP overtaking the other and preserves the low frequent changes in the wind field. In addition to the weight-parameters used in [75] for the wind speed, a new set of weight parameters are introduced for the wind direction state. This allows for a more adaptive tuning. As discussed in Section 5.3, we assume the wind direction to have a much larger area of "validity" than the wind speed, which we assume to be more local. Additionally, we also assume that wind direction measurements change more uniformly, which is why we decrease the weight with time stronger than with the wind speed estimates. To implement this change, the matrix $\mathbf{A}_{L,Wf}(\mathbf{x}_{WF})$ in Equation (5.1) has to be extended by a weighting matrix \mathbf{W}_{WF} , as well as weighting its input:

$$\tilde{\mathbf{A}}_{L,Wf}(k, \mathbf{x}_L, \mathbf{x}_{WF}) = \mathbf{A}_{L,Wf}(\mathbf{W}_{WF}(k, \mathbf{x}_L) \mathbf{x}_{WF}) \mathbf{W}_{WF}(k, \mathbf{x}_L). \quad (5.7)$$

The structure of $\mathbf{W}_{WF}(k, \mathbf{x}_L)$ is given in Appendix 5.A.2. The changed version of Equation (5.1) is given in Equation (5.8). A visual representation of the weights can be seen as part of the results in Figure 5.11.

5.2.2. ENSEMBLE KALMAN FILTER FORMULATION

The Ensemble Kalman Filter works by employing a model multiple times with varying starting states [111]. Additionally, it is assumed, as with all Kalman Filter formulations, that both the states and the measurements are corrupted with a Gaussian white noise disturbance. This randomly generated noise, along with the different starting states, ensures that the states of the different ensembles diverge over time. The average of a state across all ensembles is its estimate and the variance its uncertainty.

Equation (5.1) is adapted as follows:

$$\underbrace{\begin{bmatrix} \mathbf{x}_L(k+1) \\ \mathbf{x}_T(k+1) \\ \hat{\mathbf{x}}_{WF}(k+1) \end{bmatrix}}_{\mathbf{x}(k+1)} = \underbrace{\begin{bmatrix} \mathbf{A}_{L,L} & \mathbf{0} & \tilde{\mathbf{A}}_{L,WF}(k, \mathbf{x}_L, \hat{\mathbf{x}}_{WF}) \\ \mathbf{0} & \mathbf{A}_{T,T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{WF,WF} \end{bmatrix}}_{\mathbf{A}_1} \underbrace{\begin{bmatrix} \mathbf{x}_L(k) \\ \mathbf{x}_T(k) \\ \hat{\mathbf{x}}_{WF}(k) \end{bmatrix}}_{\mathbf{x}(k)} + \dots \\
 \underbrace{\begin{bmatrix} \delta(\mathbf{x}_L(k), \mathbf{x}_T(k), \hat{\mathbf{x}}_{WF}(k)) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}}_{\mathbf{A}_2} + \underbrace{\begin{bmatrix} \mathbf{B}_L & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_{WF} \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} \mathbf{I}_T(k) \\ \mathbf{x}_{T,0}(k) \\ \hat{\mathbf{x}}_{WF,0}(k) \end{bmatrix}}_{\mathbf{u}} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \boldsymbol{\mu}_{WF} \end{bmatrix}}_{\boldsymbol{\mu}}, \quad (5.8)$$

where $\boldsymbol{\mu}_{WF}$ is the added noise to the wind field state. The noise combines the noise for wind speed $\boldsymbol{\mu}_u \sim \mathcal{N}(0, \mathbf{Q}_u)$ and wind direction $\boldsymbol{\mu}_\phi \sim \mathcal{N}(0, \mathbf{Q}_\phi)$, which are assumed to be Gaussian noise. All ensembles are propagated in time using Equation (5.8) with individually generated noise.

5

Due to the fact that $\hat{\mathbf{x}}_{WF}$ is perturbed individually for all ensembles, and that the values of $\hat{\mathbf{x}}_{WF}$ are coupled to \mathbf{x}_L , this state also changes differently for all ensembles. Depending on how far the ensembles have diverged from one another, the $\hat{\mathbf{x}}_{WF}$ states are at different locations in the different ensembles. The EnKF framework however assumes the states to describe the same location. This is an inherent characteristic for all simulations which include, or consist of, Lagrangian markers, particles traveling based on their own state. In literature there are two proposed ways to address this issue: The first option is to map the ensemble states to a common grid. The state correction is then applied at common locations. The corrected states are then mapped back onto the individual ensemble states. An alternative is to enhance the state with the markers position. These are then also estimated and corrected. The idea is that the location of the OPs is correlated with the states causing the propagation. This method can significantly increase the size of the problem, but also returns more information. In the case of FLORIDyn, this framework would then also correct the location of the wake. To decrease the complexity of the problem, we apply the former method. Rather than creating a new grid, we first calculate the mean position of the OPs $\bar{\mathbf{x}}_L$ across all ensembles and then apply the weighting (Equation (5.6)) to find the representative state of the ensemble at $\bar{\mathbf{x}}_L$. This can be seen as a coordinate transformation from the states of the ensemble to the mean states of all ensembles. The corrected state would then need to be projected back onto the ensemble states, however, an inversion of the weighting matrix is numerically difficult. We therefore further simplify this step by assuming that the inverse is equal to an identity matrix. This assumption is supported by the fact that the weighting matrix has a diagonally dominant structure for the OPs which are in the wind field area. It is to be expected that this assumption can not hold for more diverging wind directions and for areas where no state correction is possible.

From this point onwards we will refer to the resulting state of Equation 5.8 as forecast state, marked by an f as upper left index. If measurements are available, the forecast state $^f\hat{\mathbf{x}}_{WF, e_i}$ of ensemble e_i is corrected using the difference in predicted outputs and measured outputs with Equation (5.9). The result is the analysis state $^a\hat{\mathbf{x}}_{WF, e_i}$. The general formulation of the analysis step is as follows:

$${}^a\hat{\mathbf{x}}_{WF, e_i} = f\hat{\mathbf{x}}_{WF, e_i} + \mathbf{K} [\mathbf{d}_{e_i} - \mathbf{g}(\mathbf{x}_{L, e_i}, \mathbf{x}_T, f\hat{\mathbf{x}}_{WF, e_i})] \quad (5.9)$$

$${}^a\hat{\mathbf{x}}_{WF, e_i} = f\hat{\mathbf{x}}_{WF, e_i}, \quad (5.10)$$

where \mathbf{K} is the Kalman gain matrix, $\mathbf{g}(\mathbf{x}_{L, e_i}, \mathbf{x}_T, f\hat{\mathbf{x}}_{WF, e_i})$ describes the nonlinear output function which converts the ensemble field state to the predicted measurements and \mathbf{d}_{e_i} is a set of polluted system measurements. We assume \mathbf{x}_T to be equal across all ensembles. If no measurements are available, Equation (5.10) is used to determine the analysis state instead of Equation (5.9).

We assume that the wind direction and wind speed are uncorrelated and can be corrected independently. Therefore, the calculation of the Kalman gain matrix is split into \mathbf{K}_u to correct the wind speed and \mathbf{K}_ϕ to correct the wind direction. For the correction of the wind speed the power generated is used as a nonlinear output, for the wind direction we assume a direct measurement at the turbine location. Therefore, the output function also varies and \mathbf{K}_u and \mathbf{K}_ϕ have to be calculated in a similar, yet different manner. This will be discussed in Section 5.2.2 and 5.2.2 respectively. Figure 5.2 shows a block diagram of the correction from the point of view of a single ensemble. It visualizes how the estimated wind field states are converted into estimates at the turbine location. The wind direction is determined by a time varying output matrix, the wind speed by a nonlinear function which converts the wind speed in power generated. Both estimates are then compared to the polluted measurements. The difference is multiplied with the respective Kalman gain matrix and the forecast states are corrected and fed back into the ensemble. Figure 5.3 visualizes which calculations are executed for all ensembles, and which calculations only once. The calculations are mainly based on the EnKF literature and adapted for this case, for more information on the EnKF see [111, 124].

5

ENKF CORRECTION OF THE WIND SPEED

To calculate the Kalman gain matrix \mathbf{K}_u we follow an approach also used in [117] which allows the use of a nonlinear output function. First the averaged state error matrix $\mathbf{E}_{f\hat{\mathbf{x}}_{WF, u}}$ is calculated:

$$\mathbf{E}_{f\hat{\mathbf{x}}_{WF, u}} = \frac{1}{\sqrt{n_e - 1}} [f\hat{\mathbf{x}}_{WF, u, e_1} - f\bar{\mathbf{x}}_{WF, u}, f\hat{\mathbf{x}}_{WF, u, e_2} - f\bar{\mathbf{x}}_{WF, u}, \dots, f\hat{\mathbf{x}}_{WF, u, e_{n_e}} - f\bar{\mathbf{x}}_{WF, u}], \quad (5.11)$$

where $f\hat{\mathbf{x}}_{WF, u, e_i}$ describes the wind speed estimate of the i -th ensemble and $f\bar{\mathbf{x}}_{WF, u}$ describes the average across all n_e ensembles. The same is done for the output of the ensembles:

$$\mathbf{E}_{\hat{\mathbf{p}}} = \frac{1}{\sqrt{n_e - 1}} [\hat{\mathbf{p}}_{e_1} - \bar{\mathbf{p}}, \hat{\mathbf{p}}_{e_2} - \bar{\mathbf{p}}, \dots, \hat{\mathbf{p}}_{e_{n_e}} - \bar{\mathbf{p}}], \quad (5.12)$$

where $\hat{\mathbf{p}}_{e_i}$ denotes the estimation of power generated by the i -th ensemble, which is directly dependent on $f\hat{\mathbf{x}}_{WF, u}$ as shown in Equation (5.2) and (5.3). We can now calculate the state-output error correlation covariance matrix, Equation (5.13), and the output error covariance matrix, Equation (5.14):




Figure 5.2.: Block diagram of an ensemble with correction. The left box (—) describes the state propagation within the ensemble. If no measurements are available, the system keeps progressing the state and all ensembles diverge due to the influence of noise. The right box (---) describes the comparison of the measurements with the estimates and the resulting correction. Elements connected to the wind direction are colored in dark blue, elements related to the wind speed are colored in orange. Products and functions have a one-line frame, values have a double line frame.

$${}^e\mathcal{C}_{f_{\hat{\mathbf{x}}_{WF,u}}}, \hat{\mathbf{p}} = \mathbf{E}_{f_{\hat{\mathbf{x}}_{WF,u}}} \mathbf{E}_{\hat{\mathbf{p}}}^T \quad (5.13)$$

$${}^e\mathcal{C}_{\hat{\mathbf{p}}, \hat{\mathbf{p}}} = \mathbf{E}_{\hat{\mathbf{p}}} \mathbf{E}_{\hat{\mathbf{p}}}^T, \quad (5.14)$$

In parallel, the power measurements \mathbf{P} from the wind turbines have to be polluted:

$$\mathbf{P}_{e_i} = \mathbf{P} + \epsilon_{i,P} \quad \epsilon_{i,P} \sim \mathcal{N}(0, \mathbf{R}_P) \quad (5.15)$$

where $\epsilon_{i,P}$ is an artificial error with a Gaussian distribution, an average of 0 and the covariance matrix \mathbf{R}_P . The matrix \mathbf{R}_P is a parameter which needs to be set based on prior knowledge and tuning. If no random perturbations are added to the measurements, the variance of the analyzed ensembles becomes too low [125]. The wind speed is then corrected as follows:

$$\mathbf{K}_u = {}^e\mathcal{C}_{f_{\hat{\mathbf{x}}_{WF,u}}}, \hat{\mathbf{p}} [{}^e\mathcal{C}_{\hat{\mathbf{p}}, \hat{\mathbf{p}}} + \mathbf{R}_P]^{-1} \quad (5.16)$$

$${}^a\hat{\mathbf{x}}_{WF,u,e_i} = f_{\hat{\mathbf{x}}_{WF,u,e_i}} + \mathbf{K}_u [\mathbf{P}_{e_i} - \hat{\mathbf{p}}_{e_i}(f_{\hat{\mathbf{x}}_{WF,u,e_i}})] . \quad (5.17)$$

ENKF CORRECTION OF THE WIND DIRECTION

We assume to have a measurement of the wind direction available at the turbine locations. The state error ${}^f\mathbf{E}_{\mathbf{x}_{WF,\varphi}}$ is calculated equivalently to Equation (5.11). Contrary to the wind speed estimation, it is then used to calculate the state error covariance matrix:

$${}^e\mathbf{C}_{\mathbf{x}_{WF,\varphi}, \mathbf{x}_{WF,\varphi}} = {}^f\mathbf{E}_{\mathbf{x}_{WF,\varphi}} {}^f\mathbf{E}_{\mathbf{x}_{WF,\varphi}}^T. \quad (5.18)$$

To get the relation between the states and the output, the output matrix \mathbf{C}_{φ,e_i} is needed. It is given by the rows of the weighting matrix $\mathbf{W}_{WF}(k, \mathbf{x}_L)$ which combine ${}^f\mathbf{x}_{WF,\varphi}$ to a wind direction estimate at the OP at the rotor plane, see Equation (5.7). Due to the fact that OPs are slightly differently located in each ensemble, \mathbf{C}_{φ,e_i} is slightly different in each ensemble. We assume that $\mathbf{C}_{\varphi,e_1} \approx \mathbf{C}_{\varphi,e_2} \dots \mathbf{C}_{\varphi,e_{n_e}}$ to a degree where the basic assumptions of the EnKF still hold. However, we are diverging from the traditional calculation by generating individual \mathbf{K}_{φ,e_i} for each ensemble, based on the different \mathbf{C}_{φ,e_i} matrices. The wind direction measurements φ are also polluted with an error $\epsilon_{i,\varphi} \sim \mathcal{N}(0, \mathbf{R}_\varphi)$, equal to the power measurements in Equation (5.15). As with \mathbf{R}_p , \mathbf{R}_φ is a parameter which needs to be set. The resulting analysis step is described by:

$$\mathbf{K}_{\varphi,e_i} = {}^e\mathbf{C}_{\mathbf{x}_{WF,\varphi}, \mathbf{x}_{WF,\varphi}} \mathbf{C}_{\varphi,e_i}^T \left[\mathbf{C}_{\varphi,e_i} {}^e\mathbf{C}_{\mathbf{x}_{WF,\varphi}, \mathbf{x}_{WF,\varphi}} \mathbf{C}_{\varphi,e_i}^T + \mathbf{R}_\varphi \right]^{-1} \quad (5.19)$$

$${}^a\hat{\mathbf{x}}_{WF,\varphi,e_i} = {}^f\hat{\mathbf{x}}_{WF,\varphi,e_i} + \mathbf{K}_{\varphi,e_i} \left[\varphi_{e_i} - \mathbf{C}_{\varphi,e_i} {}^f\hat{\mathbf{x}}_{WF,\varphi,e_i} \right]. \quad (5.20)$$

Note how the Equation pair (5.19) & (5.20) differs from (5.16) & (5.17): Equation (5.19) uses the state-error-covariance matrix and a linear output matrix, where Equation (5.16) uses the output-error-covariance matrix and the output-to-state-error-covariance matrix. As a result, Equation (5.20) corrects based on a linear relation of the output to the estimated state, while Equation (5.17) compares outputs with a nonlinear relation to the estimated state.

Figure 5.3 shows the difference between the single \mathbf{K}_u for the wind speed reduction and the multiple matrices to correct the wind direction. Generally, this approach is not desirable as it requires more calculations and therefore more computational effort. We chose it as we had access to the exact output matrices and were therefore able to reduce the number of approximations in the correction. Furthermore, the later discussed test case is a nine turbine case, which means that the to-be-inverted part of the Kalman gain calculation is a 9×9 matrix, which resembles a manageable computational cost. As the number of turbine grows, a single \mathbf{K}_u matrix becomes much more desirable.

LOCALIZATION

The EnKF works under the premise that enough ensembles are simulated to approximate the correlation among the states and measurements. In order to decrease the number of ensembles needed for the error covariance approximation, prior knowledge of the system can be used to modify the covariance matrices. In practice, this is done by calculating an additional covariance matrix based on the distance of the states to each other. This covariance matrix is then multiplied with the state-error covariance matrix and the output-to-state error covariance matrix to generate localized versions without nonphysical cross-correlations:




Figure 5.3.: Block Diagram of the correction across all ensembles. The yellow, top boxes describe all ensembles, running in parallel. If no measurements are given, the forecast state is used as analysis state and the ensembles diverge further. If measurements are given, the wind speed (orange) and direction (dark blue) states are corrected. The measurements and their perturbations are given at the bottom of the figure with thin double lined frames. Without perturbation, the variance of the ensembles would become too low [125].

$$\begin{aligned}
 {}^e \tilde{C}_{\mathbf{x}_{WF}, \varphi, \mathbf{x}_{WF}, \varphi} &= C_{\mathbf{x}_L, \mathbf{x}_L} \circ {}^e C_{\mathbf{x}_{WF}, \varphi, \mathbf{x}_{WF}, \varphi} \\
 {}^e \tilde{C}_{f_{\hat{\mathbf{x}}_{WF}, u}, \hat{\mathbf{p}}} &= C_{\mathbf{x}_L, \mathbf{x}_{L,T}} \circ {}^e C_{f_{\hat{\mathbf{x}}_{WF}, u}, \hat{\mathbf{p}}} \\
 {}^e \tilde{C}_{\hat{\mathbf{p}}, \hat{\mathbf{p}}} &= C_{\mathbf{x}_{L,T}, \mathbf{x}_{L,T}} \circ {}^e C_{\hat{\mathbf{p}}, \hat{\mathbf{p}}},
 \end{aligned} \tag{5.21}$$

where \circ is the Hadamard product, $C_{\mathbf{x}_L, \mathbf{x}_L}$ is the localization covariance between all states, $C_{\mathbf{x}_L, \mathbf{x}_{L,T}}$ the localization covariance between all states and the OPs at the turbine locations and $C_{\mathbf{x}_{L,T}, \mathbf{x}_{L,T}}$ the localization covariance between the OPs at the turbine locations among each other. As the initial OPs are placed at the turbine location, $C_{\mathbf{x}_{L,T}, \mathbf{x}_{L,T}}$ and $C_{\mathbf{x}_L, \mathbf{x}_{L,T}}$ are subsets of $C_{\mathbf{x}_L, \mathbf{x}_L}$. To calculate the localization covariance we follow [126] and use a piece wise defined polynomial with a characteristic cut-off length l . Every element $c_{i,j}(d, l)$ of the covariance matrix connecting OP_i and OP_j is then calculated as follows, based on their distance d to one another:

$$c_{i,j}(d, l) = \begin{cases} 1 & d = 0 \\ -\frac{1}{4} \left[\frac{d}{l} \right]^5 + \frac{1}{2} \left[\frac{d}{l} \right]^4 + \frac{5}{8} \left[\frac{d}{l} \right]^3 - \frac{5}{3} \left[\frac{d}{l} \right]^2 + 1 & 0 < d \leq l \\ -\frac{1}{12} \left[\frac{d}{l} \right]^5 - \frac{1}{2} \left[\frac{d}{l} \right]^4 - \frac{5}{8} \left[\frac{d}{l} \right]^3 - \frac{5}{3} \left[\frac{d}{l} \right]^2 - \frac{5}{7} \frac{d}{l} + 4 - \frac{2}{3} \left[\frac{d}{l} \right]^{-1} & l < d \leq 2l \\ 0 & 2l < d. \end{cases} \tag{5.22}$$

The equation is based on the work of [127], which also offers alternative functions. The state and output covariance matrices in Equation (5.16) and (5.19) are subsequently replaced by their modified equivalents, defined in Equation (5.21).

5.3. RESULTS AND DISCUSSION

This section presents the results of the introduced framework. The parameters chosen for the simulation are discussed in Section 5.3.1. In Section 5.3.2, we apply the framework to a test case where FLORIDyn is used as both the reference system and the ensemble system. In a second step, Section 5.3.3 showcases a setup where an LES simulation is used as reference.

5.3.1. ENSEMBLE KALMAN FILTER AND LOCALIZATION PARAMETERS

Table 5.1 lists the used parameters for the EnKF framework and the new weighting method introduced in FLORIDyn. For the other FLORIDyn parameters, see [47]. The number of ensembles defines how well the covariance matrices are approximated, but with every ensemble also a new FLORIDyn simulation has to run. Preliminary tests suggested that 50 ensembles yields acceptable results: the estimate becomes noisy for few ensembles (< 20) and the results do not change significantly for higher numbers (tested up to 150). The correction frequency C_f determines how often the framework gets called and has to be a multiple of the FLORIDyn simulation frequency, which is set to 0.25 s^{-1} . The noise on the measurements and states is assumed to be uncorrelated and is therefore described by an identity matrix, multiplied by a factor. The factors were chosen based on the variance of the measurements in the LES simulation. The noise perturbation for the wind direction had to be increased to allow for a faster correction. The cut-off length for the localization function was set in accordance to be roughly double of the spatial component of the weighted average function, scaled by $\sqrt{10/3}$, which is motivated by the findings of [128]. The values for the weighted average calculation of the wind speed are inspired by [75] but were modified. The initial values were lower for $\sigma_{w,t,u}$ and higher for $\sigma_{w,dw,u}$. This led to unreasonable weighting areas during major wind direction changes. Therefore, we adapted the values to hold value longer but over shorter downwind distance. The weight for the crosswind distance $\sigma_{w,cw,u}$ was also set lower, which led to very little interaction between downstream turbines and passing wakes. The value was therefore increased. The wind direction weights are set by manual tuning and intuition: as the wind direction changes, the entire flow field changes and we can assume a more global effect than with wind speed. If the age of the OPs would not be punished more heavily by the weighting, old wind direction measurements would prohibit a change. We therefore chose to give the wind direction state a wide influence area but a short lifespan. All values should be seen as educated estimates and not necessarily as part of a final parameter set.

5.3.2. FLORIDYN AS VALIDATION PLATFORM

To understand the behavior of the developed algorithm, we first use FLORIDyn itself as a validation platform. This allows us to model flow conditions, which could be hard to generate in high-fidelity simulations. We also can carve out the differences of the FLORIDyn model by itself and FLORIDyn within the EnKF

Table 5.1.: Parameters used in the Ensemble Kalman Filter framework

Ensemble Kalman Filter		Localization	
n_e (-)	C_f (s^{-1})	l_u (m)	l_ϕ (m)
50	1/12	$\sqrt{10/3} \cdot 500$	$\sqrt{10/3} \cdot 1000$
Noise per 4 s			
R_P (MW)	R_ϕ (deg)	Q_u (ms^{-1})	Q_ϕ (deg)
$0.1 \cdot I_{n_T}$	$3 \cdot I_{n_T}$	$0.4 \cdot I_{n_{OP}}$	$3 \cdot I_{n_{OP}}$
Weighted Average			
$\sigma_{w,dw,u}$ (m)	$\sigma_{w,cw,u}$ (m)	$\sigma_{w,t,u}$ (s)	$\sigma_{w,dw,\phi}$ (m)
256	126	256	512
			$\sigma_{w,cw,\phi}$ (m)
			512
			$\sigma_{w,t,\phi}$ (s)
			50




Figure 5.4.: The left figure depicts the wind farm layout which is used in Section 5.3.2 and 5.3.3, along with the main wind directions in both cases. The right figure depicts the transition between the flow variables for the different turbines during the FLORIDyn reference simulation.

framework. This gives us an idea about the contribution of the EnKF framework. To this end, we consider a 9-turbine case subject to heterogeneous wind speed and wind direction changes. The 3×3 wind farm layout is described in Figure 5.4. All nine DTU 10 MW turbines [33] are placed with equidistant spacing of 900 m, which is roughly equivalent to 5 turbine diameters. The FLORIDyn model is propagated as described in Equation (5.8), excluding the noise term. During the 1200 s simulation duration, both the wind direction and speed change heterogeneously throughout the field. The wind direction is initialized with 255 deg and then changes at rates between 0.2 deg s^{-1} and 0.05 deg s^{-1} to 280 deg. Figure 5.4 shows the start and end directions, as well as the transitions for the different turbines. The wind speed is changed in a similar manner from 8 m s^{-1} to 10 m s^{-1} at rates between 0.02 m s^{-2} and 0.01 m s^{-2} . Yaw-misalignment is not part of the presented case.

Figure 5.5 depicts the power generated by the center row turbines, T3, T4 and T5. The leftmost turbine, T3, experiences only free stream conditions and the only change in power generated stems from the increase of wind speed. As the EnKF needs to adapt to the changing wind speed, the predicted power generated trails the reference with a delay of $\approx 10 \text{ s}$. The other two turbines, T4 and T5, are subject to upstream wakes and show a reduction in power generated due to




Figure 5.5.: FLORIDyn as reference: the black line shows the FLORIDyn simulation, the orange line the estimate with 1, 2 and 3 standard deviations.

passing wakes during the wind direction change. Both also show that the EnKF predicts the reduction in power generated, but contrary to the prior case, and counter intuitively, the EnKF does not follow the reference simulation but leads it. To understand this effect, we have to look at the state correction: the EnKF has the capability to change the states of shed OPs whereas the FLORIDyn simulation only adds new data to the OPs at the turbine locations. As a result, downstream OPs in the EnKF simulation correct their state while they initially remain unaffected in the FLORIDyn simulation. The effect is that the wake in the EnKF adapts sooner to the new wind direction and crosses the downstream turbines earlier than in the FLORIDyn simulation. This difference can be seen in Figure 5.6 where the

5




Figure 5.6.: Estimated background flow field wind direction compared to the reference at $t = 564$ s.

reference state is compared to the EnKF estimate along with the respective wake locations in white and black. The EnKF simulation shows a similar heterogeneous state distribution as the reference but at a higher value. The resulting difference shows that the EnKF is uniformly at the same or higher value than the reference simulation. The difference plot also shows how the EnKF wake has progressed further south than in the reference case. Figure 5.7 depicts a similar behavior where the reference states show bigger differences in the background wind speed than the estimated state. The estimated state is more uniform and suggests a more steady change of the wind speed throughout the flow field. To conclude, the framework creates smoother, low-pass filtered flow fields. The EnKF adds a spatial connection to the FLORIDyn simulation that it, in this implementation, would not have otherwise. This is due to the correction capabilities of the EnKF: the state-




Figure 5.7.: Estimated background flow field wind speed compared to the reference at $t = 804$ s.

error-correlation calculation and the wider localisation area connect and change states which are not necessarily connected by the spatio-temporal averaging alone, which is applied in FLORIDyn. This could be changed by narrowing the localisation window or by widening the spatio-temporal average influence. The EnKF is furthermore able to track heterogeneous flow field changes.

5

5.3.3. SOWFA AS VALIDATION PLATFORM

We now consider a more realistic and complex setting, in which we use the developed EnKF FLORIDyn framework to estimate the flow field state during a 60 deg wind direction change. We use a high-fidelity large-eddy simulation performed with SOWFA (Simulator fOr Wind Farm Applications)[129] as ground state. Section 5.3.3 describes the simulation setup and case. The wind speed and direction state estimation results are discussed in Section 5.3.3 and 5.3.3 respectively. Section 5.3.3 highlights the influence of spatial weighting on the state correction. Finally, Section 5.3.3 looks at the power generated.

SIMULATION CASE

The wind farm layout is identical to the case discussed in the previous Section 5.3.2 and is shown in Figure 5.4. Differences are that the wind direction changes from 255 deg before $t = 600$ s to 195 deg at $t = 900$ s and afterwards. In SOWFA the turbines are subject to turbulent inflow at roughly 6 % ambient turbulence intensity and an average wind speed of 8.2 m s^{-1} . The case has been used in previous FLORIDyn publications, such as [47] and has been described in greater detail in [1]. The SOWFA setup files and output data are available at [44].

For the evaluation two snapshots are considered: one at $t_1 = 600$ s, when the wakes are fully developed within the wind farm boundaries, and one at $t_2 = 840$ s, during the wind direction change. Animations of the full simulation are available at [130], along with data to recreate the flow field plots and power generated.

WAKE LOCATION AND ESTIMATED WIND SPEED

Figure 5.8 shows the SOWFA flow field at t_1 and t_2 , as well as the SOWFA flow field overlaid with the FLORIDyn OPs and contour lines. The estimated flow field overlaps well with the SOWFA simulation and follows the curved wakes, caused by the wind direction change. At its current development stage FLORIDyn can

not model wake meandering effects, which are present in the SOWFA flow field. The estimated FLORIDyn wind speed state is depicted in Figure 5.9, along with




Figure 5.8.: The left figure pair depicts the SOWFA flow field before and during the wind direction change at $t_1 = 600$ s and $t_2 = 840$ s respectively. The right pair superimposes contour lines of the estimated FLORIDyn flow field, as well as the OPs.

the standard deviation, calculated by the relation between the ensembles. The wind speed estimate is relatively uniform, which is also the case in the SOWFA simulation. Coarser patterns of lower and higher wind speeds can be seen for the three turbine rows. The lower row of T6, T7 and T8 does also show the same lower wind speeds in the SOWFA simulation (Consider Figure 5.4 as reference for the turbine numbering). The standard deviation is below 0.3 m s^{-1} and does not show a meaningful pattern. During the wind direction change, the standard deviation rises for some downstream areas but remains bounded. The filter properties of spatio-temporal averaging strongly contribute to this result, as small scale changes are averaged out. The area influenced by the averaging is shown in Figure 5.11 and is further discussed in detail in Section 5.3.3.




Figure 5.9.: The estimated effective wind speed flow field is depicted in the left figure pair, once in steady state at $t_1 = 600$ s and once at $t_2 = 840$ s, during the wind direction change. The right figure pair shows the respective standard deviation.

WIND DIRECTION STATES

The wind direction changes uniformly throughout the wind farm and is not subject to turbulent changes, such as the wind speed. Therefore, noise is only introduced by means of system noise and added measurement noise. If these are chosen too low, the EnKF trusts the model state too much and adapts too slow during the

actual state change. Knowledge about the flow field effects is crucial to set the noise magnitude: in this simulation, the wind direction changes with 0.2 deg s^{-1} . Consequently the system noise and measurement noise have to be chosen higher. The Gaussian noise is set to 0.75 deg s^{-1} for both the measurement and the system noise. In steady state however, the difference between estimate and true state is minimal. This can be seen in Figure 5.10. During the wind direction change,




Figure 5.10.: The left figure pair shows the difference to the true value and the value extrapolated from the OPs' wind direction states. A blue color indicates that the extrapolated wind direction is trailing the true value during the direction change. The true values are 255 deg at $t_1 = 600 \text{ s}$ and 207 deg at $t_2 = 840 \text{ s}$, left and right respectively. The right figure pair shows the standard deviation of the state values, based on the variance across the ensembles.

the error increases, but due to the introduced noise, the states are corrected and the error remains within small bounds within the wind farm. Outside of the wind farm little to no corrections are made and the states keep diverging across all ensembles, visible in Figure 5.11. This is also mirrored in the standard deviation plot in Figure 5.10, where the areas between the wind turbines, with wake overlap, show the lowest values.

WEIGHTING AND CORRECTIONS

The previous results underlined the importance of the applied spatio-temporal weighting, as well as the localisation. The results also show how the weighting changes the way FLORIDyn acts as a simulation and how it defines boundaries to what the EnKF can correct. This section looks at a part of the simulation in detail to quantify and understand these implications.

Figure 5.11 focuses on the upper third of the wind farm area and the states within it. During steady state the focus lies on the OP states of the turbines T0, T1 and T2. During the wind direction change the wakes of T0, T1 and T2 leave the observed area, which is why the now more present wakes of T3, T4 and T5 are discussed instead. The displayed metrics from top to bottom are the wind speed state of the OPs, the wind direction state of the OPs, the weights for the wind speed averaging, the weights for the wind direction averaging and the SOWFA flow field overlapped with the FLORIDyn contour lines.

The wind speed state plots show the estimate of all OP states for three turbines with $\pm 1, 2, 3$ standard deviations, as well as a reference line at 8.2 m s^{-1} . The reference is the mean wind speed across the wake free flow field and only gives an indicator of where the estimate should be. The figures also show the development of the standard deviation separately as black line. The leftmost states are closest




Figure 5.11.: This figure compares the simulation state at the start and during the wind direction change. The top two figure rows visualise the states of all OPs connected to Turbine 0 to 2, at $t_1 = 600$ s and of all OPs connected to Turbine 3 to 5, at $t_2 = 840$ s. The figures show the state value of the OPs with 1, 2 and 3 standard deviations (std) as orange plot and y-axis, and the std alone as black plot and y-axis. A grey area marks the OPs which have left the wind farm boundaries. The x-axis denotes the state index, e.g. 10 relates to the 10th OP. The dotted lines mark a reference value: for the wind direction, this is the true value, for the wind speed, this is the average precursor speed. Below the OP states are three rows of flow field plots. They show a third of the wind farm to allow for a more detailed look. The upper two show the weights used for the weighted average calculation, see Equation (5.6). The lower row shows the SOWFA flow field overlapped with the contour lines of the FLORIDyn simulation.

to the turbine, the rightmost furthest away. The grey areas indicate where the states leave the wind farm boundaries. Every 10th state (or OP) is marked in the weight plots to aid the interpretation across the plots.

Looking at the wind speed state it is notable that the standard deviation remains rather constant at or below 0.2 ms^{-1} within the wind farm bounds, but then

increases as the OPs move farther away from the wind farm region. This means that within the wind farm, the turbines and their measurements keep the variance bounded. As the OPs leave the wind farm bounds, this effect diminishes and the system noise leads to a divergence of the ensembles. A noticeable difference between the variance of the wind speed states and the later discussed wind direction states is the fact that the wind speed variance does not significantly decrease within the wind farm bounds. This might be expected as the OPs travel past downstream turbines, which act as correcting sensors. The lack of significant correction can be explained by the wind speed weight plot. It shows the influence by the OPs based on the spatio-temporal weighting. Due to the weighting parameters chosen for the wind speed, the area of the influence of an OP has very little overlap with a neighbouring wake. The states within a wake therefore tend to develop more independently from other wakes, in this wind direction. During the wind direction change however, the wakes cross and areas of locally lower variance can be seen in the wakes of T3 and T4. These OPs with lower variance are roughly in the direct neighbourhood of the wind turbines T1 and T2. For instance, the OPs of T3 passed T2 60 s (= 15 time steps) prior to the snapshot - this reduced the variance of the OPs in the proximity of T2, which then traveled further. The lowered variance can now be seen around OP 70 in the wake of T3.

The reduction of the variance is more dominant with the wind direction states, as the framework assumes a wide area of influence of OPs due to the spatio-temporal weighting. This allows the EnKF to cross-correct from one wake to another. The weighting plot also shows how the wide area of influence decreases with the age of the OPs: OPs close to their turbine have high values, but already the weights of OP 10 and 20 have decreased significantly. In contrast, the wind speed state weights remain longer at a high level. The wind direction weights therefore encourage corrections of OPs close to the turbines. Thus, the estimated variance decreases in the proximity of the turbines, an effect which is also visible in Fig. 5.9. The weights and therefore the connection of measurement-to-state decreases rapidly outside of the wind farm boundaries and the ensembles diverge. During the wind direction change, the states within the wind farm are successfully corrected, while states outside of the wind farm are not. Different from the wind speed states, the wind direction reference resembles the true state, which is why the value is given in absolute difference rather than absolute values.

POWER GENERATED

The power generated is not one of the estimated states, but it is inherently linked to them, as shown in Equation (5.2). Figure 5.12 depicts the power for six out of the nine turbines in the EnKF framework and SOWFA. The EnKF estimates the power generated 2700 times during the simulation, 300 values per turbine, every 4 s. Out of these, 73.5% are within one standard deviation (as estimated by the EnKF) from the SOWFA value, 90.7% within two and 95% within three. This does not quite match the Gaussian ideal distribution of 68%, 95% and 99.7%. The mean error lies at 0.0452 MW, about 1 % of the absolute power, the root-mean squared error at 0.4936 MW, which is about 11 %. The small mean error seems to justify the assumption that the surrogate model is error-free on average, which is a necessary assumption for the EnKF. However, the wind speed state noise might be underpredicted. Another metric which is sometimes used for the EnKF is




Figure 5.12.: Power generated from six of the nine turbines. This reduction was made as the other turbines are mostly in free stream conditions and do not add more information. The black line depicts the SOWFA simulation as reference, orange shows the EnKF estimate including 1, 2 and 3 standard deviation boundaries. As the wind direction changes from $t = 600$ s to 900 s, T1, T2, T4 and T5 experience a reduction in wind speed due to the passing wakes. This leads to the sudden reduction in the generated power.

5

how often the estimate is above or below the measured value. If we understand the true system as another ensemble, the ratio should be about 50%. For the power generated, this value lies at 53.56% in favor of an underestimate. Note that these simulations have been run without parameter tuning of the underlying analytical wake model, which would influence these results.

Figure 5.13 shows a comparison of the same EnKF setup with different correction times, once where the state is corrected every 12 s and once every 60 s. The vertical lines mark the times at which the states are corrected. Note, that the time window has been reduced to the last 600 s of the simulation to allow for a clearer picture.

The simulation with the larger correction step shows a much wider variance cone than the simulation with a shorter step, as the ensembles have more time to diverge. Nevertheless, the EnKF is able to track the reference signal under most conditions. A significant decrease of the estimation quality becomes apparent during the wake passing, which is much delayed, compared to the SOWFA simulation. In a dynamic wind farm control scenario, this could result in delayed actions which try to improve on a situation that already happened. The delay stems from the fact that FLORIDyn as surrogate model assumes a steady state input - during the estimation the wind direction and speed do get corrected but are then held until the next correction. This could be exchanged for a momentum driven approach, which continues to change the state based on the previous rate of change. This is subject to further research. If such an extension would be valuable can only be judged based on real measurement data and conditions. The implication for this case are that the estimated wakes adapt

too slow and the difference leads to a delayed wake overlap prediction, as well as an overshooting of the power signal. The overshooting is also present in some of the power estimates for turbines in the case with a shorter correction step, see T4 or T5 in Figure 5.12 in a similar context. The dataset [130] contains additional plots for intermediate correction times of 24, 36 and 48 s.




Figure 5.13.: Comparison of the predicted power generated by Turbine 1 in an EnKF framework which is updated every 12 s to one which is updated every 60 s. The vertical lines indicate the frequency of correction.




Figure 5.14.: Comparison of the power generated by T1 in different FLORIDyn implementations. The grey area indicates the time window during which the wind direction changes.

Lastly, Figure 5.14 compares the EnKF estimate of the power generated to previous FLORIDyn results in [47]. Previous work used the *Immersion and Invariance* (I & I) estimator to estimate the effective wind speed based on the rotor speed and generator torque [52]. The base model uses a set wind speed and a prescribed wind direction change. While the base model is unable to mirror the influence of the turbulent wind speed in the reference simulation, it is able to predict the timing of the wake overlap, as well as the approximate magnitude. Adding the I & I estimator couples FLORIDyn to the reference simulation and removes the need for a prescribed wind speed. It also allows a close tracking of the power generated. This model however still needs a prescribed wind direction. There is furthermore no state correction. The EnKF framework includes both, wind speed and direction.

It is further able to correct previously generated states and provides certainty bounds for the estimate. On the downside is that this model needs to simulate multiple versions of the same simulation, while previous results were obtained with one simulation.

5.4. CONCLUSION

The presented work formulates an ensemble-based wind farm flow field estimation framework which can estimate the background wind speed as well as the wind direction. It is based on the dynamic parametric wind farm model FLORIDyn and utilizes the Ensemble Kalman Filter approach to correct wind speed and direction across the simulation. The approach was tested in a 3×3 wind farm case with heterogeneous and changing flow conditions. The results show, that the framework is able to follow the flow field changes. Its estimate and corrections strongly depend on the spatio-temporal averaging and localization parameters. These can limit but also enable state corrections and are a vital part of the inner workings of framework. This work also shows that the EnKF can be used to estimate the state of purely Lagrangian particle based simulations. It further gives an idea on how a state augmentation with the particle position can be avoided by instead projecting the ensembles onto common particle positions. This is under the assumption that the particles across different ensembles are not too far apart.

Given the broader context of wind farm flow control, this work presents an essential building block for a realistic closed-loop dynamic control approach for operational wind farms. Future work will need to address how the framework works in larger wind farms and what the correct parameter choices are under realistic circumstances, as well as what framework extensions are necessary.

5.A. STATE SPACE DESCRIPTION

This appendix describes two aspects of the state space description of FLORIDyn more in detail: Appendix 5.A.1 presents the system matrix and the input matrix of the FLORIDyn model. Appendix 5.A.2 then discusses how spatio-temporal weighting is applied in the system matrix.

5.A.1. SIMPLE PROPAGATION

The following equations describe the state propagation in the FLORIDyn model and the structure of the matrices:

$$\begin{bmatrix} \mathbf{x}_L(k+1) \\ \mathbf{x}_T(k+1) \\ \mathbf{x}_{WF}(k+1) \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{L,L} & \mathbf{0} & \mathbf{A}_{L,WF}(\mathbf{x}_{WF}(k)) \\ \mathbf{0} & \mathbf{A}_{T,T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{WF,WF} \end{bmatrix} \begin{bmatrix} \mathbf{x}_L(k) \\ \mathbf{x}_T(k) \\ \mathbf{x}_{WF}(k) \end{bmatrix} + \dots \\ + \begin{bmatrix} \boldsymbol{\delta}(\mathbf{x}_L(k), \mathbf{x}_T(k), \mathbf{x}_{WF}(k)) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{B}_L & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_{WF} \end{bmatrix} \begin{bmatrix} \mathbf{I}_T(k) \\ \mathbf{x}_{T,0}(k) \\ \mathbf{x}_{WF,0}(k) \end{bmatrix}. \quad (5.23)$$

The matrix $\mathbf{A}_{L,L}$ is given by

$$\mathbf{A}_{L,L} = \begin{bmatrix} \mathbf{A}_{L,L,T_1} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{L,L,T_2} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{A}_{L,L,T_{n_T-1}} & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \mathbf{A}_{L,L,T_{n_T}} \end{bmatrix}_{n_T \cdot n_{OP,T} \cdot n_L \times n_T \cdot n_{OP,T} \cdot n_L} \quad (5.24)$$

$$\mathbf{A}_{L,L,T_i} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \dots & \dots & \mathbf{0} \\ \mathbf{A}_{L,L,OP_1} & \mathbf{0} & \dots & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{L,L,OP_2} & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{A}_{L,L,OP_{n_{OP,T}}} & \mathbf{0} \end{bmatrix}_{n_{OP,T} \cdot n_L \times n_{OP,T} \cdot n_L} \quad (5.25)$$

$$\mathbf{A}_{L,L,OP_i} = \mathbf{I}_{n_L}, \quad (5.26)$$

where \mathbf{I}_{n_L} is the $n_L \times n_L$ identity matrix and n_L describes the number of location states per OP. The number of turbines is given by n_T and $n_{OP,T}$ describes the number of OPs per turbine. The matrices $\mathbf{A}_{T,T}$ and $\mathbf{A}_{WF,WF}$ are similar, they only differ in the size of their smallest components: the matrices $\mathbf{A}_{T,T,OP}$ and $\mathbf{A}_{WF,WF,OP}$ differ in size as the number of stored turbine and wind field states is different.

The matrix $\mathbf{A}_{L,WF}(\mathbf{x}_{WF}(k))$ is described by the same structure than $\mathbf{A}_{L,L}$, only the smallest component differs:

$$\mathbf{A}_{L,WF,OP} = \begin{bmatrix} \Delta t \cos(\varphi_{OP}) & \mathbf{0} \\ \Delta t \sin(\varphi_{OP}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}_{n_L \times n_{WF}} \quad (5.27)$$

where Δt is the time step of the simulation and φ_{OP} is the wind direction at the location of the OP. The number of wind field states is given by n_{WF} . We assume here that the first state is the wind speed.

The input matrices feed inputs into the first OPs of the turbine. For the location data, this is defined as follows:

$$\mathbf{B}_L = \begin{bmatrix} \mathbf{B}_{L,T_1} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_{L,T_2} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{B}_{L,T_{n_T-1}} & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{B}_{L,T_{n_T}} \end{bmatrix}_{n_T \cdot n_L \cdot n_{OP,T} \times n_T \cdot n_L} \quad (5.28)$$

$$\mathbf{B}_{L,T_i} = \begin{bmatrix} \mathbf{I}_3 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}_{n_L \cdot n_{OP,T} \times n_L} \quad (5.29)$$

We assume here the same use of location states as described in [1, 47], where world coordinates are stacked on wake coordinates. The other two input matrices are defined similarly, with the difference that \mathbf{B}_{T,T_i} and \mathbf{B}_{WF,T_i} consist only out of an identity matrix in the first rows and not as in \mathbf{B}_{L,T_i} accompanied by zero columns:

$$\mathbf{B}_{WF,T_i} = \begin{bmatrix} \mathbf{I}_{n_{WF}} \\ \mathbf{0} \end{bmatrix}_{n_{WF} \cdot n_{OP,T} \times n_{WF}} \quad (5.30)$$

5.A.2. WEIGHTED PROPAGATION

To apply a spatio-temporal weighting of the wind field states the time-varying matrix $\mathbf{W}_{WF}(k, \mathbf{x}_L)$ is introduced and applied as follows:

$$\tilde{\mathbf{A}}_{L,WF} = \mathbf{A}_{L,WF}(\mathbf{W}_{WF}(k, \mathbf{x}_L) \mathbf{x}_{WF}) \mathbf{W}_{WF}(k, \mathbf{x}_L). \quad (5.31)$$

The matrix weights all OP states with respect to a location of interest. The location of interest in this case are the OPs themselves. As a result, row i contains the weights to sum the influence of all OP states onto OP i . The weights are calculated by Equation (5.6) and then normalized by the sum of the row. To calculate the time weight, the location of the OP within the state vector can be used along with the time step of the simulation Δt . If a variable time step is used, this needs to be adapted. In the presented work, only the wind speed and direction are weighted. Weights are denoted as $w_{i,j,u}$ or $w_{i,j,\varphi}$, where i is the OP of interest, j the OP influencing OP $_i$ and u or φ represent the wind speed or direction respectively. Furthermore, we use three wind field states: The wind speed, direction and ambient turbulence intensity. No weighting is applied to the latter, but could be implemented in a similar fashion. The matrix \mathbf{W}_{WF} is then assembled as follows:


$$\mathbf{W}_{WF}(k, \mathbf{x}_L) = \dots$$

$$\begin{bmatrix} \mathbf{W}_{WF,1,1}(k, \mathbf{x}_L) & \mathbf{W}_{WF,1,2}(k, \mathbf{x}_L) & \dots & \mathbf{W}_{WF,1,n_{OP}}(k, \mathbf{x}_L) \\ \mathbf{W}_{WF,2,1}(k, \mathbf{x}_L) & \mathbf{W}_{WF,2,2}(k, \mathbf{x}_L) & \dots & \mathbf{W}_{WF,2,n_{OP}}(k, \mathbf{x}_L) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{W}_{WF,n_{OP},1}(k, \mathbf{x}_L) & \mathbf{W}_{WF,n_{OP},2}(k, \mathbf{x}_L) & \dots & \mathbf{W}_{WF,n_{OP},n_{OP}}(k, \mathbf{x}_L) \end{bmatrix} \quad (5.32)$$

$$\mathbf{W}_{WF,i,j}(k, \mathbf{x}_L) = \begin{bmatrix} w_{i,j,u}(k, \mathbf{x}_L) & 0 & 0 \\ 0 & w_{i,j,\varphi}(k, \mathbf{x}_L) & 0 \\ 0 & 0 & w_{i,j,l_0} \end{bmatrix} \quad (5.33)$$

$$w_{i,j,l_0} = \begin{cases} 1 & \text{if } i=j \\ 0 & \text{otherwise} \end{cases} \quad (5.34)$$

Note that the row-sum of $\mathbf{W}_{WF}(k, \mathbf{x}_L)$ must return a vector of $\mathbf{1}$.



6

TIME-SHIFTED COST FUNCTION DESIGN

This chapter is based on the following publication:

[131] M. Becker, D. Allaerts, and J.W. van Wingerden, *Time-shifted cost function design for more efficient dynamic wind farm flow control*, 2024 IEEE Conference on Control Technology and Applications (CCTA) (2024), 10.1109/CCTA60707.2024.10666535.

Dynamic wind farm flow control is the art and science to maximize the energy yield of large wind farms. In this paper we will address the problem of large time delays between control actions of the different turbines in the farm and the delayed impact on the downstream turbines. We propose and show how a time-shifted cost function approach can render the receding horizon optimization problem more efficient and can mitigate the unavoidable turn-pike effect. We further show how the resulting setup can be used to break the optimization problem apart into several smaller optimization tasks to reduce the computational load. We demonstrate that the proposed changes do allow an economic model predictive control strategy to engage into collaborative wind farm control for long term gains, while a more traditional cost function approach leads to greedy turbine behavior. As a result, we take a crucial step towards a mature implementation of dynamic model based wind farm flow control.

6.1. INTRODUCTION

Wind farms form an integral part in the renewable energy mix of the future. They combine the necessary infrastructure with an allocated location for kinetic wind power extraction. In this context a conflict arises: The more wind turbines are placed in the same area, the more they influence one another. As one turbine extracts the kinetic energy from the wind, it leaves a wake of lower wind speed behind it. This leads to power losses at turbines operating in the wake of an upstream turbine [132].

The goal of wind farm flow control is to mitigate these wake losses by changing the turbine state and by extend the wake shape. One such control strategy is wake steering, where the rotor plane is intentionally misaligned with the main wind direction [50, 133]. As a result, the wake is deflected and can be steered away from downstream turbines. Steady-state control approaches do exist for this strategy, and have been found effective in field experiments [4]. These approaches do neglect the wake dynamics by design, so the question remains if a dynamic model-based control approach can improve upon the results of a steady state control strategy. While dynamic control-oriented wind farm control models have been proposed [1, 134] there are limited to no results available on how these models can outperform existing control algorithms.

One central barrier for dynamic wind farm flow control are the time delays between the control actions and the effect on downstream turbines. The delays make it necessary to simulate the wind farm at least until the point at which the upstream control changes arrive at the downstream turbine and often more time is needed to escape the turnpike effect [109, 135]. The turnpike effect arises when optimization parameters have effects which lie beyond the prediction horizon of the simulation. In the case of dynamic wind farm flow control, this usually leads to a greedy control of the upstream turbine: As it needs to sacrifice some of its power to enable the gains of a downstream turbine, it chooses to maximise its own power again at the end of the time horizon. The resulting disadvantages at the downstream turbine lie beyond the time horizon and are not captured by the cost function anymore. In practice this means that a part of the solution found by the optimization problem has to be disregarded. Workarounds do exist: One way to reduce the impact is to fixate the last variables of the action horizon [95], or to attach large costs to moving them. Faced with a similar issue, [84] uses a time shifted signal to identify the impact of a dithering signal in a model-free control approach.

We propose to extend the approach of using a time shifted signal, and to dynamically choose which turbines to consider for the cost function based on the time it takes for control actions to propagate downstream. This way, the cost function is reduced to the impact of the control actions at the upstream turbine synchronised with the impact at the downstream turbine(s).

The main contributions of this paper are threefold:

- We propose a new time-shifted cost function for dynamic wind farm control to mitigate the turnpike effect.
- We propose a novel clustering algorithm to break down the control problem in smaller uncoupled optimization problems.
- We demonstrate the concept with a proof-of-concept simulation study.

This paper is structured as follows: Section 6.2 describes the methodology used. In Section 6.3, the main results are presented. Finally, in Section 6.4 the conclusions will be drawn.

6.2. METHODOLOGY

We first describe the general problem setup in Sec. 6.2.1. In Sec. 6.2.2 we derive the time delays between the turbines. Using these delays we deduce a way to optimize the turbine behaviour across a given action horizon in Sec. 6.2.3. Sec. 6.2.4 comments on the choice of the prediction and action horizon length and on means to further reduce the computational effort.

6.2.1. GENERAL SETUP

The receding horizon problem as we use it consists of an action horizon of τ_{ah} time steps, and a prediction horizon of τ_{ph} time steps. Fig. 6.1 (a) depicts this time line. For each time step in the action horizon and for each turbine we consider a optimization variable $\theta_{i,k}$, where $i \in [1, n_T]$ is the index of a turbine (e.g. T_1) and $k \in [1, \tau_{ah}]$ is a time step. In our problem, these optimization variables directly relate to the turbine yaw angle γ in degrees. Therefore, they are also constrained: The absolute difference between two consecutive parameters may not be larger than the turbine yaw rate limit allows. In Fig. 6.1 (a) this is indicated by the triangular areas behind the trajectory. Fig. 6.1 (b) shows a sketch of a possible reaction of the power generated by a two turbine wind farm, T_1 and T_2 , to the yaw angle of T_1 : T_1 initially sacrifices power by changing its yaw angle. Meanwhile T_2 produces very little power as it is negatively affected by the wake of T_1 . As the change propagates through the wake and arrives at T_2 , the downstream turbine is less impacted by the wake and the power increases.

6.2.2. TIME DELAY ESTIMATION

In order to determine when a downstream turbine will be affected by the control signal, we need to approximate the advection speed. This is the speed with which the changes in the flow field propagate downstream. Two predominant approaches exist to determine the advection speed: (i) the use of a constant fraction of the free wind speed or (ii) the use of a wake model and the resulting knowledge of the effective wind speed in the wake. The latter method is used for instance by [75] and comes at the computational cost of solving the related model equations. The approach of a constant fraction of the free wind speed has been numerically derived and used by [84] and [32], and has been experimentally studied by [136]. In this work we assume $u_{adv} = 0.7 \cdot u_{free}$, following [84] as they are focusing on the same turbine-to-turbine interactions relevant to control. The time offset between two turbines is therefore calculated by

$$t_{i \rightarrow j} = \frac{x_{T_i \rightarrow T_j}}{0.7 \cdot u_{free}}, \quad (6.1)$$

where $t_{i \rightarrow j}$ describes the time it takes for changes at the rotor plane of T_i to arrive at T_j . The downwind distance from T_i to T_j is denoted by $x_{T_i \rightarrow T_j}$. Since we consider




Figure 6.1.: A receding horizon example where the optimization parameters θ_{T1} of turbine T1, here representative for the yaw angle γ_{T1} , are adapted in (a) from the current time step until τ_{ah} . The effect of these changes can be observed over τ_{ph} time steps in (b). The power of T1 suffers due to the yaw angle change, but the power of T2 increases after some delay. The wind turbines and wakes are sketched in (c) .

discrete time steps, we round $t_{i \rightarrow j}$ to the closest previous time step:

$$\tau_{i \rightarrow j} = \left\lfloor \frac{t_{i \rightarrow j}}{\Delta t} \right\rfloor. \quad (6.2)$$

Equally important to the downstream distance of the turbine is the crossstream distance due to the wake expansion. In this work we disregard turbines with a larger crossstream distance than $\pm 2 D$, where D stands for the turbine diameter.

6.2.3. COST FUNCTION ASSEMBLY

Our goal is to move from a single cost function to multiple smaller cost functions: To this end we initially consider all optimization parameters to be independent. Then, if two or more parameters affect the same system output, they become part of a set. For each set of optimization parameters we derive one cost function that combines their impact on the system outputs. This process is described in greater detail in the following paragraphs.




Figure 6.2.: (a) The relational matrix T^* between the turbine outputs P (columns) and the optimization parameters θ (rows). The combination of linked optimization parameters leads to (b), indicated by one example. Combining consecutive outputs then leads to T , depicted in (c), with one highlighted example.

Basic formulation Each optimization parameter $\theta_{i,k}$ belongs to a turbine T_i with $i \in [1, n_T]$ and a time step $k \in [1, \tau_{ah}]$. The output of the turbine, its power, is

denoted as $P_{T_i}(k)$. As time marches, each $\theta_{i,k}$ has an initial effect on $P_{T_i}(k)$, but then delayed with $\tau_{i \rightarrow j}$ time steps also on the downstream turbine T_j and thus $P_{T_j}(k + \tau_{i \rightarrow j})$. To bring this relation into matrix form we define the vectors \mathbf{P} and $\boldsymbol{\theta}$:

$$\mathbf{P} = \begin{bmatrix} P_{T_1}(1) \\ P_{T_1}(2) \\ \vdots \\ P_{T_{n_T}}(\tau_{ph}) \end{bmatrix}, \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_{1,1} \\ \theta_{1,2} \\ \vdots \\ \theta_{n_T, \tau_{ah}} \end{bmatrix} \quad (6.3)$$

We then define a matrix \mathbf{T}^* . The (i, j) -th element of \mathbf{T}^* is 1 if the i -th element of $\boldsymbol{\theta}$ affects the j -th element of \mathbf{P} and 0 otherwise. Fig. 6.2(a) depicts an example how this matrix could look like for a four turbine wind farm.

Combine coupled optimization parameters If one column of \mathbf{T}^* has two or more 1-entries, multiple optimization parameters affect this output. This means that we cannot consider these optimization parameters as independent and need to optimize them together. We combine the entire rows of the affected parameters by a logical OR operation and combine the parameters into a new set \mathcal{O}_i . Fig. 6.2(b) shows which rows of (a) have been combined.

Connect consecutive sets The previous step combined all parameters that need to be connected. The following step is not required but further reduces the number of optimization parameter sets for convenience: If two sets, \mathcal{O}_i and \mathcal{O}_j , optimize consecutive-in-time parameter sets they can be combined: $\mathcal{O}_i^* = \mathcal{O}_i \cup \mathcal{O}_j$, where \mathcal{O}_i^* denotes the updated set. In practice this means that consecutive time steps are connected to a single, longer time series. This does change how many optimization parameters are in one set, but it does not change how many turbines are considered in one optimization. The condition for the corresponding operation in the relational matrix is a diagonal similarity of the rows. See Fig. 6.2(b) where diagonally equal rows share the same color. Their combination leads to the matrix \mathbf{T} , depicted in Fig. 6.2(c).

Cost function derivation Each row of \mathbf{T} describes a way to map \mathbf{P} to a number of smaller optimization problems that each optimize a subset of $\boldsymbol{\theta}$. In our case, this is the relation between the power generated by the turbines and the yaw angle of the turbines. The resulting cost function is the negative sum of the power generated:

$$\min_{\boldsymbol{\theta}_i} J_i(\boldsymbol{\theta}_i) = -\mathbf{t}_i \mathbf{P}(\boldsymbol{\theta}_i) \quad (6.4)$$

where i relates to optimization problem i , $\boldsymbol{\theta}_i$ denotes the optimization parameters in \mathcal{O}_i and \mathbf{t}_i is the i -th row of \mathbf{T} . In this particular case the cost of actuation has an intrinsic negative impact on the power, which is why we do not include it again as a negative term in the cost function.

Simulation order & clusters Eq. (6.4) is a general description of all optimization problems to solve. What it does not offer is an order in which the problems need to be solved. For our system, we have to ensure that the yaw rate is below a given limit. This means that we depend on an initial yaw angle at the start of our

optimization problem in order to deduct the consecutive ones. This dependency can be derived from \mathbf{T} . Fig 6.3 shows as an example how the optimization problem in the first row relies on previously deducted results for T_2 , T_3 , and T_4 in row 2. Row 2 further relies on row 3 and row 3 on row 4. Therefore, the solve-order of the optimization problems is $\theta_4 \rightarrow \theta_3 \rightarrow \theta_2 \rightarrow \theta_1$. Two things are worth noting: first, to solve optimization 4 we only need to simulate two time steps. This means that we can reduce the number of optimization parameters, the number of simulated turbines and also the time. Secondly, in a larger wind farms independent solving graphs appear, as not all turbines do influence each other. This offers to reduce the computational cost by splitting the optimization in smaller, independent wind farm clusters.¹




Figure 6.3.: The arrows indicate which row depends on the results of another row. In this case row 1 depends on 2, which depends on 3, which depends on 4. As a result, we need to solve optimization problem $4 \rightarrow 3 \rightarrow 2 \rightarrow 1$. The blocks indicate which part relates to which turbine. In order to solve row 4, we only need to evaluate T_4 .

6

6.2.4. CHOOSING τ_{AH} AND τ_{PH}

One of the central motivations of this work was to avoid the turnpike effect. The time shifted cost function approach introduced in Sec. 6.2.3 does a first step towards that goal by only combining turbine outputs that have an effect on the optimization parameters. But a poor choice of the prediction and action horizon will lead to a turnpike effect nevertheless. For instance, if we chose $\tau_{ph} = \tau_{ah} = 1$ the effect of the control action will never reach a downstream turbine and the turbines will behave greedy. To avoid this we use the results from Sec. 6.2.2 and set a minimum for τ_{ph} :

$$\tau_{ph} \geq \tau_{ah} + \max_{i,j \in [1, n_T]} \{ \tau_{T_i \rightarrow T_j} \}. \quad (6.5)$$

It may be valid to set a $\tau_{ph, \max}$ which violates Eq. (6.5). The entries of the diagonal block in \mathbf{T} are then shifted beyond τ_{ph} and are neglected. The violation could be based on an assumed maximum wake length or to lower the computational cost. Another measure to reduce computational cost would be to not simulate turbines at times when their output is not used, see the zero-columns in Fig. 6.2(a-d). In that case we alter Eq. (6.5) to only simulate the turbine j for as long as needed by other turbines:

$$\tau_{ph,j} \geq \tau_{ah} + \max_{i \in [1, n_T]} \{ \tau_{T_i \rightarrow T_j} \}. \quad (6.6)$$

Note that it might still be necessary to simulate the wakes of the turbines. A third variation would be to define individual action horizons for all turbines, based on a

¹The Matlab code to generate \mathbf{T} and other derivatives is publicly available with the DOI 10.4121/54cfbca7-243e-4a27-af2d-74cdd91471b2

given τ_{ph} :

$$\tau_{ah,i} \leq \tau_{ph} - \max_{j \in [1, n_\tau]} \{ \tau_{T_i \rightarrow T_j} \}. \quad (6.7)$$

This can be especially advantageous for systems where changes propagate through the domain, e.g. in the case of a wind speed change or gust. A front row turbine might have a short action horizon, because it can only react to changes in the wind, while downstream turbines can take the wind history of the front row turbine as prediction into account.

6.3. RESULTS

In this section we describe the simulation case along with the wind farm flow control setup in Sec. 6.3.1. The results of the case are presented and discussed in Sec. 6.3.2.

6.3.1. SIMULATION

We use the *Flow Redirection and Induction Dynamics* (FLORIDyn) framework [47] as eMPC model to evaluate the cost function and to carry out the reference simulation. The FLORIDyn framework has been designed to simulate wind turbine wake dynamics at a low computational cost, and is aimed to be used in a model based closed-loop wind farm flow control strategy. To test the proposed cost function design, we carry out three simulations:

1. a baseline simulation in which the turbines perfectly track the wind direction to maximise their own power generated,
2. a “naive” economic model predictive control (eMPC) setup where FLORIDyn is used to optimize the power generated over a given prediction horizon,
3. an eMPC setup with the same optimization setup but with the proposed cost function structure.

The difference between approach 2) and 3) isolates the effect of the proposed cost function.

We use a three DTU 10-MW reference turbine [33] wind farm, depicted in Fig. 6.4. The synthetic wind direction signal is set to maintain a steady wind direction of 260 deg for the first 1000 s, followed by a constant change to 290 deg over the following 1000 s. The 290 deg are maintained for the future. This setup does encourage both, individual power maximisation of unwaked turbines and collaborative wind farm flow control to maximise the wind farm performance. In addition, the wind direction change forces changes in the yaw set points of the turbines. The simulation is carried out with a free wind speed of 8 ms^{-1} , an ambient wake model intrinsic turbulence intensity level of 6 %, and with shear, using the power-law and a coefficient of 0.28 [137].

The simulation is discretized in 5 s steps. We arbitrarily choose $\tau_{ah} = 6 \rightarrow 30 \text{ s}$ for all three turbines, which results in 18 optimization variables. The prediction horizon for both eMPC strategies is calculated based on Eq. (6.5). Every control iteration, the first two steps of τ_{ah} are applied and the receding horizon is pushed forward by 10 s. The optimization is constrained by a maximum yaw misalignment by $\pm 30 \text{ deg}$ and a maximum allowed rate of change of 0.3 deg s^{-1} . The optimization is solved by the interior point algorithm for constrained minimization




Figure 6.4.: The three turbine wind farm during the wind direction change. The turbines are placed with a $5 D$ distance along the West-East axis and the centre turbine is placed $-0.5 D$ on the South-North axis. The initial wind direction is 260° , which gradually changes after 1000 s to 290° .

problems [138]. Both eMPC strategies are initialized with a yaw misalignment of 0° and have knowledge of the future development of the wind direction. This assumption is made for simplicity, similar to the work in related publications, see [95, 139].

6

6.3.2. SIMULATION RESULTS AND DISCUSSION

The yaw trajectories of all three turbines are depicted in Fig. 6.5, where the yaw misalignment is calculated as the difference between the main wind direction and the turbine orientation. The baseline controller tracks the wind direction perfectly and shows a yaw misalignment of 0° at all times.

The “naive” eMPC strategy initially uses its degrees of freedom to drive the yaw misalignment to 0° . It maintains a yaw misalignment of 0° up until the wind direction changes. In the second half of the simulation, the controller engages into yaw steering. The reason is that it tries to avoid power losses in the near future: During the cost function evaluation the controller is aware that the wind direction will continue to change. However, the short action horizon only allows it to move during the initial 6 steps of the prediction horizon. As a result, the controller moves to a future wind direction to be more aligned in the near future.

The eMPC strategy with the proposed cost function does engage into yaw steering and misaligns T2 with the main wind direction. This steers the wake away from the downstream turbine T3 and allows an increase in the total power generated by the wind farm of 0.17 MW . Towards the wind direction transient, the controller preemptively steers T1 to avoid wake interaction at a later stage, which results in large power gains during the initial transient. Turbine T1 remains yawed for a longer time but then recovers its position and reduces its misalignment with the wind direction.

Fig. 6.6 depicts the generated power by each turbine over time. The baseline and “naive” eMPC strategy show almost identical behaviour, as expected from the yaw angle trajectories. However the data of the eMPC strategy with the proposed




Figure 6.5.: Yaw angle of T1-T3 during the simulation for the baseline simulation, the “naive” eMPC strategy, and eMPC strategy with the proposed cost function structure. The dotted line indicates the activation of the controller, the dashed line the start of the wind direction change.




Figure 6.6.: The power generated by the three turbines in comparison. The dotted line indicates the activation of the controller, the dashed line the start of the wind direction change.

cost function does show that the strategy creates gains at the downstream turbines: During the steady state part of the simulation, T2 elevates the power generated by T3, which is the result of the wake steering depicted in Fig. 6.5. During the initial wind direction change, the yaw steering efforts of T1 show an effect at both T2 and T3. While the power for the two other controllers reduces as the wake interaction increases, the proposed controller is able to offset the wind speed reduction. As a joint effort of T1 and T2, the power of T3 never drops as low as it does in the baseline case or with the “naive” eMPC strategy.

The wind farm efficiency during the simulation is shown in Fig. 6.7. Once the controller is activated, the proposed eMPC strategy is willing to sacrifice efficiency, and therefore power, to gain efficiency in the long term. The proposed strategy outperforms the baseline as a result, after the effects of the induced yaw angle changes have propagated downstream. As observed in the yaw angle trajectories, the “naive” eMPC strategy does not engage in yaw steering and misses the opportunity to outperform the baseline. During the wind direction transient, the proposed method manages to offset the large drop in farm efficiency, connected to the inevitable wake passing. It is here, where the proposed method shows the largest gains, at its peak 11 % over the baseline. But as the change in wind direction progresses, the wind turbines need to recover their alignment with the main wind direction, something which the baseline has already done. Here, up to 7.6 % wind farm efficiency is sacrificed compared to the baseline.




Figure 6.7.: Wind farm efficiency η_{WF} calculated by total power generated divided by power of the farm without wake effects and misalignment. The dotted line indicates the activation of the controller, the dashed line the start of the wind direction change. The plot contains wind farm start up effects prior to the controller activation.

In summary, the proposed cost function construction allows the eMPC framework to utilise the given degrees of freedom to engage into collaborative wind farm flow control for long term gains. With a more traditional cost function, the optimization result becomes greedy and only tries to minimise short term losses due to yaw

misalignment.

6.4. CONCLUSION

In this paper we present a novel method that is able to automatically restructure the optimization of a dynamic eMPC setup for wind farm flow control. It detects sub-optimization problems based on the inherent time delays in the wind farm and the spatial distance of wind turbines. The method further returns which problems can be solved in parallel and which need to be solved sequentially. Based on these factors, new cost functions are constructed which allow an eMPC framework with a short action horizon to optimize for long term wind farm gains in both steady state and dynamic conditions. The alternative, traditional, implementation of the cost function leads to greedy control behaviour, unable to perform collaborative control.

Future work should investigate the robustness of the proposed framework, mainly towards the assumptions around the advection speed of the wake and the flow preview. The solver and an efficient use of the optimization parameters are two additional aspects not discussed in this work, which are essential to a successful eMPC design for wind farm control.



7

CLOSED-LOOP MODEL-PREDICTIVE WIND FARM FLOW CONTROL UNDER TIME-VARYING INFLOW USING FLORIDYN

This chapter is based on the following publication:

[140] M. Becker, M.J. van den Broek, D. Allaerts, and J.W. van Wingerden, *Closed-loop model-predictive wind farm flow control under time-varying inflow using FLORIDyn*, accepted for publication in *Wind Energy* (2025)

Wind farm flow control has been a key research focus in recent years, driven by the idea that a collectively operating wind farm can outperform individually controlled turbines. Control strategies are predominantly applied in an open-loop manner, where the current flow conditions are used to look up precomputed steady-state set points. Closed-loop control approaches, on the other hand, take measurements from the farm into account and optimize their set points online, which makes them more flexible and resilient.

This paper introduces a closed-loop model-predictive wind farm controller using the dynamic engineering model FLORIDyn to maximize the energy generated by a ten-turbine wind farm. The framework consists of an Ensemble Kalman Filter to continuously correct the flow field estimate, as well as a novel optimization strategy. To this end the paper discusses two dynamic ways to maximize the farm energy and compares this to the current look-up table industry standard. The framework relies solely on turbine measurements without using a flow field preview. In a 3-hour case study with time-varying conditions, the derived controllers achieve an overall energy gain of 3 to 4.4 % with noise-free wind direction measurements. If disturbed and biased measurements are used, this performance decreases to 1.9 to 3 % over the greedy control baseline with the same measurements. The comparison to look-up table controllers shows that the closed-loop framework performance is more robust to disturbed measurements but can only match the performance in noise-free conditions.

7.1. INTRODUCTION

A switch away from fossil fuels to less greenhouse gas-emitting (GHG) sources of energy is necessary to prevent a climate crisis [141]. Wind energy is one alternative that provides energy at a fraction of the GHG emissions. Wind farms are, therefore, an essential part of the energy transition. However, they do not provide the maximum amount of energy they could. How come? This is in part due to the way that turbines interact: As a wind turbine converts kinetic energy from the surrounding airflow into electricity, it leaves behind an area of low wind speed called a wake. In wind farms, these wakes will likely influence downstream turbines, lowering their power output. Wind farm flow control (WFFC) strategies aim to mitigate this effect by manipulating the wake shape.

In this work, we focus on model-based approaches to WFFC. With this approach, a model of the farm is used as a surrogate for the real wind farm. The model predicts how the turbine wakes behave, given the atmospheric conditions and turbine states. Different types of models exist with varying costs and capabilities. On the one end, high-fidelity models like Large Eddy Simulations (LES) provide the most insight into the physical phenomena that take place in a wind farm, e.g. [20, 63]. These models are typically too demanding for control which has motivated model simplifications to achieve faster computational speeds. Medium-fidelity models capture a coarse image of the flow. Two-dimensional Reynolds-averaged Navier–Stokes solvers like e.g., [45, 134, 142] fall in this category. They aim to capture the core wake dynamics at a reduced cost by limiting the dimensionality of the flow. However, later research showed that this simplification can render them useless for wake steering applications [66]. Other models in the category of medium-fidelity models simulate the wake propagation based on synthetic turbulence. The Dynamic Wake Meandering (DWM) model, introduced by [17], propagates the wake as a series of turbulence boxes. The propagation speed and direction are then determined by the contents of the box. This approach, coupled with an aeroelastic turbine model, can give estimates of loads onto the turbine and its structure. Successor models like FAST.Farm [72] and HAWC2Farm [143] provide a basis to investigate damage equivalent loads on a farm scale at a relatively low computational cost. Another approach to model wind turbine wakes is to simulate free-vortex particles shed by the rotor, e.g., [68, 109]. The downside of these models is that they tend to become numerically unstable in the far-wake region and under turbulent conditions. Low-fidelity models provide a flow field prediction at a very low computational cost. This is achieved by modeling the wake shape and wind speed reduction as a set of analytical equations, e.g., [16]. These models typically predict the steady-state wake shape of a single turbine wake and use superposition methods to combine the effect of multiple wakes in a farm. Within WFFC applications, they are used to test and optimize yaw-angle set points for an entire farm.

Based on the computational speed provided by the steady-state engineering models, paired with the success of the DWM model, an additional group of models has been proposed: These models use passive Lagrangian tracers, called Observation Points, to propagate turbine and flow field states from each turbine downstream. The Flow Redirection and induction Dynamics model FLORIDyn initially proposed this approach [19]. Since its introduction, the model has since been revised [1] and further developed [47]. Similar modeling approaches exist, e.g., [74, 75]. Their simplicity, simulated wake dynamics, and speed make them

attractive for model-based WFFC applications.




Figure 7.1.: (a) Closed-loop compared to (b) open-loop wind farm flow control. The figure is adapted from [132].

Model-based WFFC is predominantly applied in an open-loop configuration [132]. This means that the yaw angle set points are optimized ahead of time; see Figure 7.1 (b). This is mainly done by employing a steady-state wake model and selecting a set of ambient conditions. For each ambient condition, the yaw angle set points are then optimized and stored in a Look-up Table (LuT). During operation, the current ambient conditions, such as wind speed and direction, are identified and used to look up the optimal yaw angles. Examples of this strategy can be found in the field experiments conducted by e.g., [4, 61, 144]. The issue with this approach is that the strategy can not react to unforeseen circumstances. Examples of this could be wind farm layout changes due to offline turbines, unmodeled or incorrectly modeled turbine interactions, or heterogeneous and changing flow conditions.

This is addressed by closing the loop, see Figure 7.1 (a), where the model is continuously updated by measurements. A recent approach is to close the loop on the parameters of the steady-state model [51, 110, 145, 146]. This entails a coupling between the measured flow conditions and the modeled ones. The difference is then fed back to correct aspects like the wake expansion. In [147], they employ a similar strategy, but rather than correcting the model parameters, they build a corrector for the model output. This version of closed-loop control still inherently neglects the wake dynamics. Closed-loop wind farm flow control using a dynamic wake model like FLORIDyn is still a poorly explored area of research. [21] employs the first version of the FLORIDyn model paired with a Kalman Filter to estimate the wind speed in a six-turbine wind farm. The same publication also uses FLORIDyn to perform a yaw angle optimization; however, not in closed-loop.

Control for energy maximization using dynamic models uses predominantly a receding horizon approach called Model-predictive control (MPC), e.g., [109, 131, 148, 149]. This means that the optimization is done over a predetermined prediction horizon. The surrogate model is used to predict how the cost function will be impacted by the control actions taken. An optimizer is then used to determine the ideal control actions. The optimization is followed by an update time step, which is typically smaller than the prediction horizon. During this time, the previously optimized time series of control set points is applied. After the update time step has passed, a new optimization is done. MPC is typically applied to follow a reference value, e.g. a reference wind farm power [81]. In contrast

this work aims to maximize a cost function, which leads to a different optimization problem. This is referred to as economic MPC (eMPC).

A re-occurring assumption in this context is the full knowledge of the flow field (e.g. [148–151]), and full knowledge of the environmental changes ahead (e.g., [95, 109]). The latter is often referred to as preview and can lead to significant gains over preview-less control approaches, also for steady-state control approaches [95, 139, 152]. Yet, it is not clear how this preview information may be attained in a realistic WFFC scenario. The same holds true for the previously mentioned assumption of full flow field knowledge. State-estimation techniques can provide some of this knowledge. They use sensor data to correct the current state of the surrogate model. This is done by comparing predicted measurements to taken ones. Looking at FLORIDyn specifically, in the work by [21], the turbine power generated is used to correct the local wind speed estimate in FLORIDyn by employing a Kalman Filter. Similarly, [83] also uses the power generated, as well as wind direction measurements, to correct a heterogeneous flow field state in FLORIDyn using an Ensemble Kalman Filter (EnKF). This methodology has the advantage that it provides the state estimate as well as an uncertainty estimate. Additionally, it does not require a linear(-ization) of the model. In the later work of [90] and [91], the wake location is corrected by using an EnKF and a downstream turbine as a coarse sensor [153].




Figure 7.2.: Closed-loop design applied in this paper. The wind farm provides power and wind direction measurements at the turbine locations, which are used to correct the FLORIDyn model state. The identified state is then used to optimize the future yaw angle set points for all turbines. This is passed on to a low-level controller, which applies the set points.

The literature review shows that there is a gap in closed-loop wind farm flow control approaches using dynamic surrogate models in time-varying conditions. Prior work has developed surrogate models like FLORIDyn to dynamically model the wakes in a farm. It has also proposed state-estimation frameworks to align the model's state with the true state. How to then use FLORIDyn to obtain the optimal control set points is still an unexplored topic.

This paper, therefore, proposes a novel closed-loop economic model predictive control framework for farm-wide energy maximization. To this end, the paper builds upon previous work with the following four main contributions: (i) a discussion around the nature of the cost functions for energy and power maximization, (ii) a dimensionality reduction of the optimization problem, which also leads to a realistic turbine operation, (iii) an implicit way to limit misalignment angles and to avoid nonlinear constraints, and (iv) the derivation of a closed-loop framework for

dynamic model-based wind farm flow control. The proposed framework focuses on readily available sensor data, like turbine power and wind direction measurements at the turbine locations, and in contrast to previous work, does not assume knowledge of the full flow field nor a preview of future flow field changes. A block diagram of the closed-loop design is depicted in Figure 7.2. The derived optimization strategy is applied in a ten-turbine closed-loop case study. This is done with a wind farm simulated in LES under turbulent conditions and with realistic time-varying wind direction changes. In this context, the paper also addresses the impact of noisy and biased measurements on controller performance and yaw travel.

The remainder of the paper is structured as follows: Section 7.2 discusses the methodology of the closed-loop approach. This consists of a description of the model (Section 7.2.1), the state estimator (Section 7.2.2), and the novel discussion related to the controller (Section 7.2.3). The simulation methods are discussed in Section 7.3. This entails the description of the high-fidelity environment, as well as the tested wind conditions, sensor data, and controller settings. The proposed framework is tested in Section 7.4, where the results are presented. Lastly, Section 7.5 draws a conclusion of the work.

7.2. METHODOLOGY

This section introduces the components of the closed-loop control approach presented in this paper: Section 7.2.1 describes the basics of the dynamic surrogate model, Section 7.2.2 of the state estimation. The main contribution of this paper, the controller, is discussed in Section 7.2.3. The list of parameters used for the components introduced in Section 7.2 can be found in Appendix 7.A, alongside the tuning approach.

7.2.1. SURROGATE MODEL

The Flow Redirection and Induction Dynamics Model (FLORIDyn) used in this study is based on the work presented in [47], with the extensions done in [83]. It simulates wake dynamics by employing so-called Observation Points (OPs), which carry states from the rotor plane downstream. The states one OP possesses are (i) turbine states (e.g., yaw angle), (ii) flow field states (e.g., wind direction), and (iii) its own positional states. The states are initialized at the rotor plane based on what the turbine measures. The OP position is advanced based on the wind speed and direction at its position.

While previously each OP would rely on its own state of the wind speed and direction to propagate, the wind speed and direction are now the result of a weighted average. This way, the OP takes neighboring OP states into account. The spatiotemporal Gaussian weighting function was first adapted by [75] for a FLORIDyn similar model and adapted in [83] to allow for better correlation between the OPs. This benefits the assumptions made with the state estimation; see Section 7.2.2. The FLORIDyn framework uses a Gaussian wake to model the wake deficit, wind speed reduction and wake deflection [16].

Unlike recently presented results (e.g. [95]) the simulations presented in this paper do not utilize any synthetic preview information of the wind direction. As previously mentioned, FLORIDyn does have flow field information stored in the OPs, which is propagated downstream. This provides other turbines with an

estimate of the flow field changes ahead. During the prediction phase, each time step, the free stream turbines copy their previous state, and downstream turbines adapt the wind speed and direction by the weighted average of the data stored in the surrounding OPs.

7.2.2. STATE ESTIMATION

The employed state estimation is based on [83] and uses an Ensemble Kalman Filter (EnKF) [124, 154]. The EnKF was chosen for three main reasons: (i) it allows the state estimation of nonlinear systems without derivation of a linearized model, (ii) it provides uncertainties of the estimated states, which can be used for robust decision making, and (iii) it has previously been used for similar applications, e.g., by [114] to estimate the flow field state of a wind farm simulation in the 2D solver WFSim [134]. Note that other state estimation formulations for FLORIDyn-like systems do exist. [21], for instance, employ an extended Kalman Filter. What distinguishes a dynamic system like FLORIDyn from other EnKF applications is that FLORIDyn is not a grid-based simulation but rather attaches its state to particles. Therefore, the flow is described by where the particles are, which may be different across the ensembles. The state estimation framework used addresses this issue by projecting the OPs of all Ensembles onto a common set, which is then corrected. To this end, two Kalman gain matrices (\mathbf{K}_u for wind speed and \mathbf{K}_ϕ for the wind direction) are derived from the common output matrices \mathbf{C}_u and \mathbf{C}_ϕ , as well as the correlation of the ensembles. This is done differently to the state estimator proposed in [83], where only the wind speed would be corrected this way, while the wind direction would be corrected with ensemble individual Kalman Gain matrices. The correction loop is depicted in Figure 7.3. The matrices \mathbf{W} define the weighted projection onto the common states, \mathbf{W}^{-1} their inverse. Since the inverse is generally not attainable, we set it to be equal to the identity matrix. Since \mathbf{W} does have a sparse and strongly diagonal shape, this approximation is valid. For more information on how the correction is derived we refer to the previously mentioned sources. There are EnKF-based designs for FLORIDyn-like models to correct the wake center, but these fall outside of the scope of this paper [90, 91]. Also outside of the scope of this paper fall online parameter estimation methods using an EnKF [51, 110].

7.2.3. CONTROLLER

Section 7.2.3 compares the differences between the steady-state cost function formulation and two dynamic cost functions. In this section, as well as in Section 7.2.3, the receding horizon optimization problem is explained, where control actions taken within an action horizon are evaluated over a prediction horizon. Section 7.2.3 to 7.2.3 further discuss measures to simplify the optimization problem solved at runtime: Section 7.2.3 investigates a basis function approach to derive the yaw angle time series, followed by a methodology to incorporate constraints into the cost function (Section 7.2.3). Section 7.2.3 discusses how the wind farm is decomposed into smaller sections to reduce the number of turbines per optimization. Lastly, Section 7.2.4 discusses the reference controllers.




Figure 7.3.: State estimation cycle for an ensemble e_i . The upper section of the figure depicts the evolution of the wind farm simulation, here depicted by $f(\cdot)$, the lower left section the extraction of the measured values from the simulation. The lower right part depicts the correction based on the mismatch of the predicted and recorded measurements. Elements with a double outline are ensemble-specific, elements with a single outline are the same for all ensembles.




Figure 7.4.: Wind farm power (a), energy (b), and shifted energy (c) efficiency of a two turbine wind farm based on the yaw control set point of T0, normalized by the case of no misalignment. Turbine T0 starts at $t = 0$ s with $\gamma_{T0} = 0$ deg and continuously changes its angle with 0.3 deg s^{-1} to the set point, where it remains. Turbine T1 is fully aligned with the wind direction. The dotted line shows the break-even line, the continuous line the ideal value of γ_{T0} based on the time. Since the shifted approach (c) moves the power signal by T1 in time, the data is shorter.

COST FUNCTION

Three cost functions are compared in this work - (i) the steady state power, (ii) the accumulated energy, and (iii) the accumulated energy disregarding the time delays between the turbines. Each function is explained in further detail below.

The steady-state approach to formulate a cost function $\mathcal{J}_{\text{steady state}}$ to maximize the farm energy is to sum the power of all n_T turbines:

$$\mathcal{J}_{\text{steady state}}(\theta) = \sum_{i=1}^{n_T} p_i(\theta), \quad (7.1)$$

where p_i is the power generated by turbine i and θ are the optimization variables. Time is implicitly taken into account, as steady-state models do not assume any delay effects. If time is taken into account, the energy is calculated as the integral of the turbine power generated over time. Equation (7.1) is extended as follows to form the dynamic cost function:

$$\mathcal{J}_{\text{dynamic}}(\theta) = \Delta t \sum_{i=1}^{n_T} \sum_{k=1}^{\tau_{\text{ph}}} p_i(\theta, k), \quad (7.2)$$

where $p_i(\theta, k)$ is the power generated by turbine i at time step k , and Δt is the time step of the simulation. The cost function now approximates the energy generated by the farm across all time steps $k \in [1, \tau_{\text{ph}}]$, where τ_{ph} is the prediction horizon of the optimization. This additional time dependence inherently changes the result for θ that we get: The power loss due to yaw control actuation needs to be recouped within the prediction horizon for the control action to be profitable. This is in strong contrast to the steady-state cost function, which assumes to have eternity available to recoup the cost of control actions.

The following example further illustrates this difference based on a two-turbine wind farm: turbine T1 is located $5D$ downstream and $0.5D$ cross stream from the upwind turbine T0, D being the turbine diameter. This configuration favors positive yaw misalignment from the first turbine for power and energy maximization. In this experiment, the control set point of T0 is varied, and the power generated from both turbines is recorded in FLORIDyn. At $t = 0$ s, T0 starts fully aligned with the wind direction and then changes its orientation with 0.3 deg s^{-1} to the reference misalignment angle. The chosen yaw speed is within the commonly used range with up to 0.5 deg s^{-1} [155].

Figure 7.4 shows the efficiency of the yaw misalignment angle set point over time, (a) based on the power generated, and (b) based on the energy generated up to the given time step. Both figures show the initial loss connected to the misalignment of turbine T0. Turbine T1 starts to experience the benefits of the misalignment after ≈ 150 to 200 s. At this point, the steady state optimum is reached for $\gamma_{T0} \approx +16$ deg, which leads to a 7 % increase in power generated compared to the baseline case. The energy generated, however, has integrated the losses leading up to the favorable wake redirection state. This means that this loss has to be recouped first before a yaw misalignment becomes a better choice than the baseline behavior. The more time is given, the more attractive larger yaw angles become: If the prediction horizon is 100 s long, $\gamma_{T0} = 0$ deg is the optimal solution to Equation (7.2), for 300 s it is $\gamma_{T0} \approx 8$ deg and for 500 s $\gamma_{T0} \approx 12$ deg is optimal. With an infinite prediction horizon, the optimal solution

of Equation (7.2) converges to the optimal solution of Equation (7.1), as the initial loss becomes increasingly negligible compared to the accumulated energy. This inherent characteristic of Equation (7.2) offers a conservative security, as it guarantees to recoup the initial investment in the given time frame under the assumption that conditions are steady. From this perspective wake steering becomes an investment-and-return problem, threatened by changing environmental conditions.

The described characteristics can also lead to unwanted behavior - if enough control degrees of freedom are given, T0 will initially engage in more aggressive control actions, followed by greedy behavior at the end of the time horizon. This is done since the downsides of the greedy control actions towards T1 fall outside of the time horizon. This is known as the "turnpike effect" and needs to be considered for dynamic wind farm flow control applications [109]. Ways to limit the turnpike effect is to not allow control actions at the end of the time horizon or to penalize them as part of the cost function.

An alternative is offered by [131]: The presented approach aims to synchronize the cause and effect of the control action with its impact on the actuated turbine as well as the downstream turbines. This formulation leads to a spatiotemporal split of the optimization problem into several smaller problems: (i) the spatial split ensures that wind turbines that do not influence one another are optimized separately, while (ii) the temporal split connects the cost function of the yaw angles with their impact. The following example illustrates the working of the algorithm using the same two-turbine wind farm. The free wind speed is set to 8ms^{-1} , and the wake advection speed is assumed to be equal to the free stream velocity. The simulation time step is $\Delta t = 5\text{ s}$. It then takes $800\text{ m}/8\text{ ms}^{-1} = 100\text{ s}$ or $\Delta k = 20$ time steps for one particle of air to reach the downstream turbine. If the time window during which a control action can be taken, the action horizon τ_{ah} , is set to $\tau_{\text{ah}} = 10$ time steps, the cost function of the yaw trajectory $\boldsymbol{\gamma}_{\text{T0}}$ of the upstream turbine T0 is formulated as follows:

$$\min \mathcal{J}_{\text{T0}}(\boldsymbol{\gamma}_{\text{T0}}) = -\Delta t \sum_{k=1}^{\tau_{\text{ah}}} p_0(\boldsymbol{\gamma}_{\text{T0}}, k) + p_1(\boldsymbol{\gamma}_{\text{T0}}, \boldsymbol{\gamma}_{\text{T1}}, k + \Delta k).$$

This formulation sums the power of T0 during its action horizon and the power of T1 once the control actions arrive at the turbine. Note that for this formulation, the two turbines are simulated from $k = 1$ until $\tau_{\text{ph},\text{T0}} = \tau_{\text{ah}} + \Delta k$, but only sections of the prediction horizon are taken into account for the cost function. The power of T1 depends on both $\boldsymbol{\gamma}_{\text{T0}}$ and $\boldsymbol{\gamma}_{\text{T1}}$. Since $\Delta k_{1 \leftarrow 0} > \tau_{\text{ah}}$, the yaw trajectory $\boldsymbol{\gamma}_{\text{T1}}$ can be optimized in a separate cost function:

$$\min \mathcal{J}_{\text{T1}}(\boldsymbol{\gamma}_{\text{T1}}) = -\Delta t \sum_{k=1}^{\tau_{\text{ah}}} p_1(\boldsymbol{\gamma}_{\text{T1}}, k).$$

For this cost function $\tau_{\text{ph},\text{T1}} = \tau_{\text{ah}}$. As a result we can formulate two optimization problems, $\mathcal{O}_1 = \min \mathcal{J}_{\text{T0}}$ and $\mathcal{O}_2 = \min \mathcal{J}_{\text{T1}}$. In this example \mathcal{O}_1 depends on $\boldsymbol{\gamma}_{\text{T1}}$, and thus on the result of \mathcal{O}_2 . Therefore, \mathcal{O}_1 cannot be solved independently. \mathcal{O}_2 is only evaluated over τ_{ah} time steps, and \mathcal{O}_1 over $\tau_{\text{ah}} + \Delta k_{1 \leftarrow 0}$ time steps, which is generally less than the otherwise necessary τ_{ph} to achieve similar yaw steering. Additionally, the smaller optimization problems can lead to better

convergence within the given optimization budget. More details about the time-shifted approach can be found in [131] or in a similar application in [84]. Figure 7.4 c) shows how the optimization landscape is affected by the change.

BASIS FUNCTION APPROACH

The optimized yaw time series must fulfill two constraints: (i) its rate of change needs to be lower or equal to the possible rate of change r_γ , limited by the actuators, and (ii) it should not exceed the prescribed maximum misalignment boundaries. This section focuses on the former aspect, while the latter is discussed in Section 7.2.3. In addition to the limited rate of change, it is desirable to derive yaw signals that can be realistically applied to existing yaw systems. Current yaw systems are comprised of a yaw drive, a bearing, and a brake [155]. To move the turbine, the brake is released and the yaw drive moves the nacelle with the maximum rate of change to the reference position. The brake is then engaged again to maintain the reference orientation. This behavior limits the set of feasible yaw time series, as, e.g., a continuously changing yaw signal is undesirable. This can be exploited to simplify the optimization problem and to reduce the number of inputs to the optimization. To this end, we rely on a basis function $\psi(\mathbf{o}, t_n)$, which takes two optimization parameters, o_1 and o_2 as arguments, as well as $t_n \in [0, 1]$, which is the normalized time within the action horizon t_{AH} . The idea is to either in- or decrease the turbine's yaw angle with the maximum rate of change and to allow the change period to start at an early or late point in time within t_{AH} . The function ψ has been designed in such a way that $o_1, o_2 \in [0, 1]$, which provides a generic interface for an optimization algorithm and can improve its effectiveness of it. The basis function is further determined by $r_{\gamma,n} = r_\gamma t_{AH}$, the maximum rate of change within the normalized parameter space.

$$\begin{aligned}\psi(\mathbf{o}, t_n) &= 2[o_1 - 0.5] \text{sat}_{[0,1]} \left(\frac{t_n - t_{s,n}(\mathbf{o})}{2|o_1 - 0.5|} \right) r_{\gamma,n}, \\ t_{s,n}(\mathbf{o}) &= o_2 [1 - 2 \cdot |o_1 - 0.5|],\end{aligned}\tag{7.3}$$

where $\text{sat}_{[0,1]}(x)$ is a saturation function which is equal to 0 for $x \leq 0$, x for $x \in [0, 1]$, and 1 otherwise. The variable o_1 determines in which direction the orientation changes and for how long, e.g. $o_1 = 1$ is a constant increase across the entire action horizon, $o_1 = 0.5$ results in a steady yaw angle. The second optimization variable, o_2 , moves the starting point of the yaw angle change period. A possible result of ψ and the resulting $\Delta\gamma$ is shown in Figure 7.5(a). The final yaw orientation trajectory is calculated as

$$\gamma(t) = \gamma(0) + \psi\left(\mathbf{o}, \frac{t}{t_{AH}}\right) \quad \forall t \in [0, t_{AH}].\tag{7.4}$$

An example is given in Figure 7.5(b). While the proposed function does limit the number of optimization parameters per turbine, it contains undesired regions of insensitivity. This namely affects the case when $o_1 = 0.5$, which results in no yaw change and removes any effect of o_2 onto the resulting trajectory. Similarly, for $o_1 = 0$ or 1 , the entire duration of t_{AH} is spent yawing, which also nullifies the impact of o_2 . This impacts the optimization landscape and introduces regions with no gradient sensitivity towards o_2 . This affects the optimizer choice in Section 7.2.3.




Figure 7.5.: Example of ψ and the resulting displacement $\Delta\gamma$ from the current orientation (a). The percentage of the action horizon τ_{ah} that is spent yawing is determined by σ_1 , while σ_2 determines the offset at which the angle change starts. At τ_u , the optimization is redone based on the updated model state. An example of how a turbine orientation might change over successive optimizations is given in (b).

IMPLICIT MISALIGNMENT LIMITATION

The methodology introduced in Section 7.2.3 limits the possible solution space of the yaw angle trajectories based on the yaw rate limit. However, it does not enforce a limit on what the resulting yaw angle might be. There are different ways of achieving a yaw misalignment limit, one would be to set it as a constraint. This is a linear constraint if the yaw misalignment is the input to the optimization problem. However, if wind direction changes are present, it is beneficial to optimize the turbine orientation instead of the misalignment, as the misalignment is a product of the uncontrollable wind direction and the turbine orientation. This switch, along with the basis function approach discussed in Section 7.2.3, leads to a nonlinear constraint of the yaw angle misalignment. This creates an additional layer to implement and can increase the complexity of the optimization problem. We instead chose to manipulate the power calculation in such a way that large yaw angles become much less attractive. This removes the constraint and makes an implicit part of the cost function. To this end, we formulate a weighting function for the power generated:

$$w_\gamma(\gamma, \gamma_{\max}, \gamma_{\min}) = \left[\frac{1}{2} \tanh(50[-\gamma + \gamma_{\max}]) + \frac{1}{2} \right] \cdots \left[-\frac{1}{2} \tanh(50[-\gamma + \gamma_{\min}]) + \frac{1}{2} \right]. \quad (7.5)$$

The weighting function reduces the power generated by 50% at the limit yaw angles and smoothly transitions from the unmodified part of the power curve to the lowered part outside of the limits. The limits were chosen as $\gamma_{\max} = -\gamma_{\min} = 33$ deg. Note that Equation (7.5) reduces the power for yaw angles that are still within but close to the bounds. As a result, yaw angles > 33 deg are already unattractive to the optimizer.

WIND FARM DECOMPOSITION




Figure 7.6.: Demonstration of the wind farm decomposition into smaller subsets, based on the $\pm 2 D$ crosswind distance selection. Turbines that are grouped together share the same color.

To reduce the computational cost of the optimization and to ensure a faster convergence, a decomposition of the wind farm in smaller subsets is useful [156]. This ensures that turbines that are unaffected by other turbines, nor affect any, can optimize for greedy behavior. Turbines that do affect one another are optimized together. In this work, we decompose the wind farm based on the current wind direction at each turbine's location. Each turbine calculates which other turbines that are within $\pm 2 D$ crosswind distance and downstream based on the current wind direction, see Figure 7.6. These then become part of the optimization group related to this turbine. We then recursively determine which turbines affect one another. This creates a directed graph, similar to the model principles proposed by [105]. Each group is then optimized together, meaning that only those turbines in the group are simulated. This effectively reduces the ten-turbine wind farm to smaller farms with one to six members. This wind farm decomposition is done with every optimization step given the current conditions and, therefore, changes over time.

OPTIMIZATION ALGORITHM

The presented framework uses a particle swarm optimization to solve the cost function, which is part of the evolutionary algorithms [157–159]. The code uses the implementation by [138]. The optimizer was chosen for the main reason that the optimization landscape is not convex. This stems from the split nature of wake steering, as steering in both directions might yield an improvement while one is favorable. Another contribution is the insensitivity of yaw angle basis function for certain parameter combinations, see Section 7.2.3. A total of 100 particles is used for all optimizations. The stopping criterion is either a maximum number of four consecutive iterations with no cost function improvement or 20 iterations. These settings were chosen after also testing 2, 10 and 40 iterations. This design further leverages the code's capability to run multiple FLORIDyn instances in parallel, similar to the EnKF. Gradient-based methods have been tested during the development of the closed-loop controllers but have been outperformed by non-gradient-based methods. The development of a dedicated optimization strategy like the Serial Refine method [93] lies outside of the scope of this paper.

7.2.4. REFERENCE CONTROLLERS

We use two types of reference dead-band lookup-table (LuT) yaw steering controllers, one that aims for 0 deg yaw misalignment and one that does implement

yaw steering. Their design is based on the work of [92] and has further been investigated in [55]. The control framework consists of two parts: (i) the dead-band filter for the wind direction and (ii) the yaw controller. The dead-band reads in the wind direction measurement $\varphi(k)$ at time step k and updates its wind speed estimate $\hat{\varphi}$ based on the difference between the two values:

$$\hat{\varphi}(k) = \begin{cases} \varphi(k) & \text{if } |\varphi(k) - \hat{\varphi}(k-1)| > \varphi_{\text{lim}}, \\ \varphi(k) & \text{if } k_i \Delta t \left| \sum_{l=\tau}^{k-1} \varphi(l) - \hat{\varphi}(k-1) \right| > \varphi_{\text{lim}}, \\ \hat{\varphi}(k-1) & \text{otherwise.} \end{cases} \quad (7.6)$$

Equation (7.6) also consists of a second update law based on the integrated difference between past measurements of φ and $\hat{\varphi}$ since the last time step τ at which $\hat{\varphi}$ was updated. The resulting $\hat{\varphi}(k)$ is then used in the LuT, denoted by f_{LuT} , which returns the yaw set points:

$$\gamma^*(k) = f_{\text{LuT}}(\hat{\varphi}(k)), \quad (7.7)$$

$$\gamma(k) = \gamma(k-1) + \text{sign}(\gamma^*(k) - \gamma(k-1)) \cdot \dots \min(|\gamma^*(k) - \gamma(k-1)|, \Delta t \delta_\gamma) \quad (7.8)$$

The turbine yaw angle $\gamma(k)$ is then updated based on Equation (7.8): It is either set to γ^* if the angle can be reached in Δt , or changes with the maximum rate of yawing δ_γ for Δt . For the baseline controller, we set $\Delta\varphi = 2$ deg and $k_i = 0.1$; for the LuT controller, we test $\Delta\varphi \in [2, 4]$ deg and $k_i = 0.1$. The LuT is calculated with FLORIDyn in steady-state conditions using the same particle swarm optimization as discussed in Section 7.2.3. This way all controllers use the same basis to make decisions. Figure 7.7 depicts the LuT for all turbines and wind directions. A characteristic of this LuT is that downstream turbines show very




Figure 7.7.: Look-up table generated with the internal FLORIDyn wake model for the reference controllers.

small misalignment angles to be optimal. This allows the downstream turbine to move slightly out of the way of a partial wake overlap. Toolboxes like FLORIS [78] avoid this behavior by setting downstream turbine misalignment angles to 0 deg by default. Since the CLC controllers tested in this work will also exhibit this behavior, it remains part of the LuT.

7.3. SIMULATION METHODS

Section 7.3.1 discusses the high fidelity model used as real surrogate, along with the wind direction case. This is followed by Section 7.3.2, which further specifies the tested controllers and used measurements.

7.3.1. “TRUE WIND FARM” SETUP

The experiments are performed in closed-loop with the LES code SOWFA as the true wind farm surrogate [20]. The wind farm is situated in a $5 \times 5 \times 1$ km domain, resolved in $300 \times 300 \times 100$ cells with an average 8 ms^{-1} free wind speed at hub height and a 0.5 s time step. All simulations are done with a neutral turbulent precursor developed for $3 \cdot 10^4$ s with a surface roughness length of $2 \cdot 10^{-4}$ m. This leads to a low turbulence intensity case and more pronounced turbine wakes compared to a high turbulence flow field. Based on early flow estimates, the TI in FLORIDyn was set to 5.4%, the precursor turbulence intensity at hub height became $I_{0, u, v, w} \approx 4\%$. The shear resulting from the surface roughness is used to calculate the power-law shear coefficient α based on the mean wind speed magnitude. The ideal shear coefficient to describe the wind speed across the rotor plane lies between 0.07 and 0.083. Based on initial precursor values, α_s was set to 0.071. The mean precursor wind speed magnitude, as well as the FLORIDyn wind speed profile, are depicted in Figure 7.8 (b). The veer across the rotor plane is less than 2 deg in the LES and is neglected in FLORIDyn.




Figure 7.8.: Precursor profile of the wind speed in the precursor and in the FLORIDyn simulation.

The closed-loop approach is tested in a three-hour long case, which is based on data collected by a vertical LiDAR at the Hollandse Kust Noord (HKN) site on the 28th of March 2023 [94]. During the measurement campaign (2019 - ongoing), wind speed and direction are recorded at different heights, with values available roughly every 20 s. First, the data was segmented into parts with sufficient data points. The values at 108 m and 133 m were used to resample and interpolate the wind direction at turbine hub height of 119 m at a regular 20 s sampling rate. The data was then low-pass filtered with a Butterworth filter with a cut-off frequency of 1/600 Hz, equivalent to [95]. The resulting sections were then investigated for completeness and interesting wind direction ramp events. For this work, one




Figure 7.9.: Schematic of the data flow for the verification and test simulations: The initial LiDAR dataset was segmented into time series of sufficient completeness, followed by resampling, interpolation, and zero-phase low-pass filtering. The stored data has then been investigated for time series of long, uninterrupted wind direction trajectories with wake interaction. Out of a 24 h period, one 3 h segment has been chosen to be simulated in a closed loop. The simulation has a precursor with the wind direction changes from the LiDAR data, which is then used to drive the simulation with the wind farm. During the simulation, measurement data is sent from the LES to a controller, which then continuously updates a table with the yaw angle set points for all turbines.

7

sub-section of 3 hours is chosen. Figure 7.9 depicts the original data and how it was prepared for the closed-loop tests. The wind farm layout and dimensions are depicted in Figure 7.10, along with the rotated case domain to simulate varying inflow conditions. The same setup has been used to conduct the LuT controller study presented in [55].

We use OpenFOAM to implement spatial uniform wind direction changes following the predefined time series of wind directions. Within the LES domain, the southwest planes are used as inflow planes, while the northeast planes are outflow planes. The farm layout and initial wind direction are rotated to fit the 225 deg inflow of the turbulent precursor. The rotated layout is centered in the domain to balance the distance to the in- and outflow planes, see Figure 7.10. The ten DTU 10 MW turbines [33] are arranged to copy a subset of the HKN wind farm and are modeled as actuator discs (ADM). Turbines modeled as ADMs on a coarser grid tend to overestimate the power generated [100, 101], which is also an issue with this setup. Therefore, with the exception of Figure 7.14, we focus on normalized power and energy quantities.

The LES environment is extended by a wind farm-wide controller that receives turbine measurements and can provide set points during the simulation runtime. The measurements received from the LES include quantities like generator power and rotor speed. Wind direction measurements are provided in one of two ways: Either directly from the data that was used to create the precursor or based on probes at hub height at the turbine locations. The former results in a noise-




Figure 7.10.: Wind farm layout of the investigated case study. The wind farm is oriented as it is in the real world, while the $5 \times 5 \times 1$ km domain is rotated to fit the inflow direction. The inflow planes are marked by wider lines. The triangle marks the wind direction range of the case. The wind direction time series is chosen such that initially T6, T8 and T9 interact, and later T0, T1, T3, T5 and T7.

free measurement of the underlying wind direction, while the latter is subject to ambient and turbine-added turbulence.

7.3.2. TESTED CONTROLLERS AND MEASUREMENTS

This section specifies the controller configurations tested in tandem with the LES and how they acquire their inputs.

Model predictive controllers All model predictive controllers use the Particle Swarm Optimizer coupled to the 2 degrees-of-freedom baseline function to derive the yaw trajectories; see Section 7.2.3 - 7.2.3. Two controllers maximize the energy over 500 s and 1000 s. This is done by evaluating the cost function in Equation (7.2). Another controller uses the shifted turbine power signals in time to synchronize control actions with their effect on downstream turbines; see Section 7.2.3. This leads to a varying prediction horizon based on the turbines involved in the current optimization problem. Typical values are between 100 s and 600 s. The controllers are summarized in Table 7.1. All controllers update the optimal yaw

Controller	Cost function	τ_{ph}
PSO, MR	Energy	$0.5 \cdot 10^4$ s
PSO, MR	Energy	$1 \cdot 10^4$ s
PSO, MR	Shifted energy	varying

Table 7.1.: Selection of closed-loop model predictive controllers. PSO refers to the Particle Swarm Optimization, MR to the maximum yawing rate basis function. All controllers have an action horizon of $\tau_{ph} = 100$ s and no preview.

set points every 60 s and use an action horizon of 100 s. This allows for a ± 30 deg orientation change based on the maximum yawing rate of $r_\gamma = 0.3$ deg s^{-1} .

Reference controllers The reference yaw-steering and baseline controllers are based on the dead-band behavior described in Section 7.2.4. The used threshold and integral gain are given in Table 7.2. In the noise-free environment the controllers are tested with a sampling time of 5 s to update the set points. With the disturbed wind direction measurements the LuT controllers are updated once every minute based on the past 1-minute average.

Steady-state model	φ_{lim}	Optimization
FLORIDyn internal	2 deg	PSO offline
FLORIDyn internal	4 deg	PSO offline
Baseline	2 deg	-

Table 7.2.: List of reference controller settings. All reference controllers have a $k_i = 0.01 \text{ s}^{-1}$

Measurements The controllers depend on measurements of wind speed and direction to provide adequate yaw angles. Table 7.3 lists the different ways data is provided to the controllers. This concerns the background wind speed u_∞ , the wind direction φ , the sampling time step, and the averaging time of the φ measurement. The noise-free data relates to the filtered LiDAR data that was

	u_∞	p	φ	Δt	average time φ
Mode 1	-	LES	noise free	15 s	0 s
Mode 2	-	LES	LES	15 s	60 s
Mode 3	given	-	noise free	5 s	0 s
Mode 4	given	-	LES	60 s	60 s

Table 7.3.: Ways for the state estimator and controller to receive data from the simulation. Mode 1 and 2 are relevant to the Ensemble Kalman Filter, Mode 3 and 4 to the reference LuT controllers.

used to drive the precursor. It can, therefore, be considered as an ideal, noise-free signal of the background flow. The same holds for u_∞ , which is provided as a constant value to the reference controllers. The EnKF for the closed-loop controllers integrates new measurements every 15 s, but a new control decision is taken every 60 s. The LuT controllers do not have a state and rather act based on the current measurement, hence the lower sampling time. Figure 7.11 showcases an example of the measurements. The black source data comes from the cleaned and zero-phase low-pass filtered field data; see Figure 7.9. The grey probe data is recorded in the LES at the rotor center of the turbine, which is, in this example, turbine T7. This was done to mimic a much-simplified version of the measurements a wind vane on a turbine might record. The plot also shows how the noise is reduced by the use of the past 1-minute averaged data instead of the raw probe data. The probe data is characterized by higher noise levels and biases during waked conditions, which poses challenges for the state estimation and the dead-band controllers. Since the wake locations are unique to every simulation, also the LuT controllers have to run online and their control actions can not be pre-computed.




Figure 7.11.: Wind direction data recorded at the location of turbine T7 during the simulation. The source data is based on low-pass filtered field measurements, which is used to drive the LES. The probe data is calculated from the u and v wind speed components at the turbine rotor center. Orange indicates the 1-minute average values of the probe data. The probe data becomes more noisy and biased as the turbine is waked by other turbines, here from 1.5 to 2.5 h. Note that the wind direction is given with respect to the LES domain.

7.4. SIMULATION RESULTS

The simulation results first investigate the wake steering behavior in Section 7.4.1. Section 7.4.2 then compares the performance of the turbines throughout the case study, followed by the farm-wide performance in Section 7.4.3.

7.4.1. WAKE STEERING

Figure 7.12 (a) depicts the orientation of turbine T4 during the last half hour of the case study. During this time, the wind direction aligns T4 with T8, as indicated by the arrow. All controllers do engage in wake steering to avoid waking turbine T8 at 10.3 D distance. However, the magnitude of the misalignment differs. Between minutes 160 and 175, the wind direction varies only marginally and does not cross the line between T4 and T8. The LuT controllers engage in the largest misalignment angles, and the Shifted CLC controller acts similarly. The CLC 1000 s controller exhibits a smaller yaw angle, and lastly, the CLC 500 s controller shows little-to-no misalignment. Figure 7.12 (b) then shows how the 10-minute energy of T8 reacts to the yaw steering efforts of T4: The more aggressive yaw angles by the LuT controllers indeed leads to a better efficiency at T8, increasing its generation by +56 %. The gain of T8, however, comes at the cost of misaligning T4 for a long time. Figure 7.12 (c) shows the combined efficiency of T4 and T8. The data shows that the period of outperforming the baseline is preceded by a period of underperformance. Based on the shallow misalignment angles, the closed-loop controllers overcome this period earlier than the LuT controllers, confirming the analysis done in Section 7.2.3 and that the CLC controllers are working as intended. They do engage in wake steering control and, based on their design, in a more or less aggressive manner. The presented example also shows that the controllers take turbines across a longer distance into account, something with which, e.g. free vortex particle models of the wake can struggle with [95].




Figure 7.12.: The orientation of turbine T4 over time as a result of the noise-free wind direction measurements is given in (a). The arrow indicates the wind direction in which turbine T8 is located at 10.3 D distance. The resulting 10-minute energy efficiency of T8 is given in (b), followed by the combined performance of T4 and T8 in (c).

7.4.2. TURBINE PERFORMANCE

Figure 7.13 shows the normalized energy gain separated for each turbine. Given the wind direction time series and the farm layout, T0, T2, T4, and T6 are upstream turbines, while T7, T8, and T9 are the most downstream turbines, see Figure 7.10.

The data shows that for all controllers, the upstream turbines sacrifice energy by yawing, which is then recouped by the downstream turbines. There are differences in the magnitude: the CLC 500 s controller amplifies the energy generation of T1, T3, and T5 more than the generation of T7, T8, and T9. It also shows a decreased investment for turbines T0, T2, T4, and T6. This is consistent with the shorted prediction horizon length and with the analysis done in Section 7.4.1. The CLC 1000 s shows more committed control actions that further lower the energy generation of the upstream turbines but also result in larger returns for the downstream turbines. The same holds true for the shifted CLC controller, as well as for the LuT controllers. An advantage that the LuT controllers have over the CLC controllers in the noise-free environment is that they are able to engage T5 and T7 consistently.

With disturbed wind direction measurements, the performance of all controllers




Figure 7.13.: Turbine individual difference between the energy generated in the controlled case and the baseline, normalized by the total baseline energy. The data in (a) relates to the controllers with noise-free wind direction measurements, and the data in (b) to the disturbed measurements. The data is normalized by the respective baseline.

decreases. The LuT controllers especially sacrifice performance with T3, T5, T7 and T9. This can be a sign that the long-distance wake interactions fail as the wind direction measurement becomes more uncertain. The CLC controllers also sacrifice performance, mainly with T1, T3, and T5. We attribute this to the lowered yaw investment by T0 due to the more uncertain wind direction.

7.4.3. FARM LEVEL PERFORMANCE

To investigate the performance of the controllers on a farm level, we compare the wind energy as the power generated over a sliding window integral of ten minutes. Figure 7.14 depicts the absolute and relative energy generated by the farm throughout the case study. The absolute data shows the magnitude at which the energy is generated. We note that the largest reductions in energy appear around the 1.5 h mark, which relates to a brief back-and-forth shift in the wind direction. This crosses the 5-turbine line T0, T1, T3, T5 and T7. During this event, the controllers return their largest gains but also losses. The losses are a product of wake steering as a method: While the turbines are slightly misaligned, the wakes are redirected to one side of the downstream turbine. If the wind direction changes, the wake can only be further redirected up to a certain point at which it is worth redirecting the wake to the other side. During this switch, the baseline wakes have likely already arrived on the new side, which means that the baseline momentarily generates more than the control strategy.

Overall, all tested controllers return consistent gains in the noise-free case, but




Figure 7.14.: Absolute (a-b) and relative (c-d) energy generated in a sliding time window of 10 minutes. The data in (a) and (c) is based on the noise-free wind direction measurements, and the data in (b) and (d) on the disturbed measurements.

7

also in the disturbed case. Here, the performance of the controllers is generally weaker due to the worse sensor data. Qualitatively, the CLC 500 s controller is the closest to the BL performance and is followed by The CLC 1000 s controller. The remaining three controllers depict a similar performance.




Figure 7.15.: Box plot of the ten-minute wind farm energy efficiency. The data is normalized with the noise-free baseline. Each pair depicts the noise-free performance on the left and the performance with disturbed wind direction measurements on the right. The median values from the left to the right are BL* 2 deg -0.5% , LuT 2 deg $+4.2\%$, LuT* 2 deg $+1.7\%$, LuT 4 deg $+4.1\%$, LuT* 4 deg $+1.4\%$, CLC 500 s $+2.2\%$, CLC* 500 s $+1.2\%$, CLC 1000 s $+2.5\%$, CLC* 1000 s $+1.5\%$, CLC shifted $+4.0\%$, CLC* shifted $+2.1\%$. The disturbed direction simulations are marked with *.

To quantify the performance, the data of Figure 7.14 (c,d) is summarized as a box plot in Figure 7.15, with the difference that all data is normalized with the noise-free baseline performance. The two LuT controllers and the time-shifted controller still depict similar performance during noise-free conditions: The LuT 2 deg and 4 deg controllers have a median efficiency of +4.2 % and +4.1 % receptively, the CLC shifted controller has a median of +4.0 %. With disturbed wind direction measurements, the performance drops to a median of +1.7 % and +1.4 % for the LuT 2 deg and 4 deg controllers, while the CLC shifted controller reduces its median to 2.1 %. We can conclude that the CLC is, therefore, more robust to sensor noise. This is mainly due to the EnKF, which by design assumes noise to be part of the sensor signal. With noise-free data, this leads to a disadvantage as the measurement is not considered to be fully “trustworthy”, but for disturbed measurements, this leads to a better estimate. This effect is also visible with the other two CLC controllers: The CLC 500 s controller drops from a median of +2.2 % to +1.0 %, the CLC 1000 s controller from +2.5 % to +1.5 %, which is a less significant decrease than the one of the LuT controllers.

A mostly neglected aspect of the cost function is the cost of actuation. It is only captured implicitly as a loss of the power generated by the actuated turbine. Actuation costs can include how much and how often the turbines yaw, how much time they spent in misalignment and how their inflow profile looks like. Recent work has suggested ways to create data-driven ways to estimate the loads on a turbine in a surrogate model [160, 161], but these are not yet included in the FLORIDyn model used in this work. We therefore resort to the yaw travel as quantity of interest.




Figure 7.16.: Energy and yaw travel of the six controllers with both noise-free and disturbed wind direction measurements. The data is normalized with the noise-free baseline controller.

Figure 7.16 shows the overall energy efficiency and normalized yaw travel. A common characteristic of the controllers under noise-free wind direction measurements is that all yaw controllers exhibit an approximately 2 to 3 × larger amount of yaw travel for a 3 to 4.5 % gain in energy. This ratio changes as the wind direction measurements become disturbed. All controllers are affected by the disturbed measurements, however at a different scale. Similarly to the trend observed with Figure 7.13, the LuT are significantly more affected by poor data

quality than the CLC controllers: All LuT controllers, including the baseline, exhibit a yaw travel amount that is ≈ 5 to $7 \times$ higher than the noise-free baseline, while the CLC controllers remain between ≈ 2.5 to $4 \times$ higher values.

7.4.4. COMPUTATIONAL PERFORMANCE

The closed-loop algorithm has two steps: (i) the state estimation and (ii) the model predictive optimization. In the case study, the LES is paused for both the state estimation and the optimization. While unrealistic, this allows for more leeway and exploration of the methodology before imposing more challenging conditions.

The state estimation is real-time compatible and roughly requires 6 s of computational time for simulation, ensemble combination, and state correction, applied every 15 s. The optimization, however, is not realtime applicable for larger numbers of turbines. Currently, groups of ≥ 3 turbines take longer than the allocated 60 s update time. These numbers are obtained in Matlab for 40 threads.

Future work should investigate the design of a dedicated optimization strategy similar to the serial refine approach [93]. And to reduce the cost per simulation. In its current form, also the FLORIDyn model used in the optimization uses the spatiotemporal weighting of the OPs. This adds a considerable cost to the model, and while it is necessary for the EnKF, it might not be for the optimization.

7.5. CONCLUSIONS

This paper introduces a novel closed-loop control framework to maximize the energy generated by a wind farm under time-varying inflow conditions based on the dynamic wake model FLORIDyn. The observed case-study results show that the framework can lead up to a median energy gain over 10 minutes of +4.0 % using a shifted cost function and +2.2 % to +2.5 % using an energy maximizing cost function. This, however, falls short behind the tested LuT controllers with gains of +4.1 % and +4.2 %. These results are obtained with noise-free wind direction measurements. If disturbed measurements from the LES are used, the performance of all controllers decreases while the yaw travel increases. However, the closed-loop controllers are less affected by this change, highlighting their robustness due to the Ensemble Kalman Filter and cost-function design. Their median 10-minute energy gain reduces to +2.1, +1.5, and +1.2 %, while the LuT performance decreases to +1.7 and +1.4 %. Future research should explore the scalability of this approach to larger wind farms, as well as in heterogeneous flow conditions.

7.A. PARAMETER TUNING

The total number of parameters is high in the proposed framework. Table 7.4 lists all parameters and constants: Six parameters to parametrize the OP weighting, one to limit the advection speed, eleven for the wake and turbine model, eight for the state estimation, and six for the optimization. Ideally, one would tune all parameters to fit a generalized scenario. The number of parameters to tune is high, which makes it difficult to replicate the related work [51, 54, 110]. We, therefore, resort to Latin Hypercube sampling of the remaining parameters and test the different combinations in a 20-minute sub-set simulation of the final setup. For each simulation, a set of farm-wide error quantities is calculated:

- Mean bias in turbine power
- Mean absolute turbine power error
- Mean squared turbine power error
- Mean squared farm power error
- Farm power bias
- Mean squared turbine power error weighted by the predicted spread of the EnKF


This is complemented by turbine individual error quantities:

- Best possible power correlation based on a variable signal time-shift
- Time-shift at which the best correlation is achieved
- Turbine power bias
- Absolute power error
- EnKF weighted power error

Out of 600 simulations, the parameter combinations that perform the best are compared for similarities. The remaining parameters are tuned manually, informed by parameters that were deemed ideal in previous studies. During this process, $\eta = 1$ was locked to further reduce the number of parameters. However, the turbine model used in the simulations overestimates the power, see Section 7.3.1. Since the power was used to correct the wind speed, the state estimation returns a $\approx 8\%$ too higher wind speed. This was captured by the error quantities and corrected by adjusting the free parameters. The wake advection factor d is smaller, and the remaining parameters lead to a faster wake recovery. The resulting parameter set may work for this setup but may not be ideal for generalization.

Parameter	Use	Range	Selected	Unit	Component
$\sigma_{w,dw,\varphi}$	weighting OP, Dir.	-	2.87	D	FLORIDyn
$\sigma_{w,cw,\varphi}$	weighting OP, Dir.	-	2.87	D	FLORIDyn
$\sigma_{w,t,\varphi}$	weighting OP, Dir.	-	50	s	FLORIDyn
$\sigma_{w,dw,u}$	weighting OP, Vel.	[0.3, 3.5]	0.6966	D	FLORIDyn
$\sigma_{w,cw,u}$	weighting OP, Vel.	[0.3, 2]	0.3570	D	FLORIDyn
$\sigma_{w,t,u}$	weighting OP, Vel.	[100, 300]	206.2331	s	FLORIDyn
d	Advection factor	[0.5, 1.0]	0.7396	-	FLORIDyn
η	Turbine efficiency	-	1.0	-	Turbine
p_p	Yaw exponent (Power)	[1.7, 2.7]	2.2	-	Turbine
α	Near wake length	-	2.32	-	Wake
β	Near wake length	[0.07, 0.39]	0.154	-	Wake
k_a	Wake expansion	[0.17, 0.92]	0.38371	-	Wake
k_b	Wake expansion	-	0.003678	-	Wake
k_{fa}	Added turbulence	-	0.73	-	Wake
k_{fb}	Added turbulence	[0, 8]	0.8325	-	Wake
k_{fc}	Added turbulence	[0, 0.5]	0.0325	-	Wake
k_{fd}	Added turbulence	-	-0.32	-	Wake
k_{TI}	TI spread	[1, 4]	3	-	Wake
$l_{loc,\varphi}$	Localisation, Dir.	-	2.8	D	EnKF
$\sigma_{\mu,\varphi}$	Process noise, Dir.	-	0*	deg	EnKF
$\sigma_{v,\varphi}$	Measurement noise, Dir.	-	3	deg	EnKF
$l_{loc,u}$	Localisation, Vel.	[3.5, 8]	6.8011	D	EnKF
$\sigma_{\mu,u}$	Process noise, Vel.	[0.1, 0.5]	0.1991	m/s	EnKF
$\sigma_{v,p}$	Measurement noise, Pow.	[0.01, 0.3]	0.08	MW	EnKF
n_e	Number of ensembles	-	50	-	EnKF
k_{enkf}	Estimation sample time	-	3	Δt	EnKF
τ_{ah}	Action horizon	-	20	Δt	eMPC
r_γ	Yaw rate limit	-	0.3	deg s ⁻¹	eMPC
$n_{iter, max}$	Max. optimization iterations	-	20	-	eMPC
k_{mpc}	Optimization sample time	-	12	Δt	eMPC
γ_{max}	Yaw limitation	-	33	deg	eMPC
γ_{min}	Yaw limitation	-	-33	deg	eMPC

Table 7.4.: Collection of all parameters and constants, their use, the investigated range, and selected value, as well as the component they belong to. The turbine diameter $D = 178.4\text{m}$ and time step $\Delta t = 5\text{ s}$ are used to normalize some parameters. Parameters without range have not been tuned. *The process noise was unintentionally set to $\sigma_{\mu,\varphi} = 0\text{ deg}$, and should, for future experiments, be set to a higher value.



8

CONCLUSIONS

In this thesis, we set out to increase the energy a wind farm generates to substitute fossil fuels further with fewer greenhouse gas-emitting alternatives. To this end, an economic model predictive closed-loop wind farm flow controller was designed. This required the development of a dynamic wake model, designing a suitable estimator, and deriving a control strategy. This has resulted in several contributions that are useful to the wind farm flow control community. The control strategy proved its capabilities in a case study under time-varying atmospheric conditions. Future work will have to assess the new controller's performance in field experiments and achieve the goal of this thesis in the real world.

8.1. CONCLUSIONS

The overarching objective of this thesis was formulated in Chapter 1 as

Thesis objective The development of a closed-loop economic model-predictive control framework using wake steering and a dynamic wake surrogate model to maximize the energy generated by a wind farm under realistic time-varying conditions.

To this end, the contributions of this thesis were split into three parts. The first part, consisting of Chapter 2 to 4, a dynamic wind farm flow control model was constructed and developed. The second part, Chapter 5, then introduced a framework to estimate the current state of the wind farm and its surrounding flow. The third part of this thesis is dedicated to the control problem of energy maximization: Chapter 6 proposed a control law based on an energy cost function that considers the delays between the turbines. This is followed by Chapter 7, which presented a case study of closed-loop wind farm flow control using different control laws and a high-fidelity model as a true wind farm. The conclusions of each part are presented below.

8.1.1. DYNAMIC WAKE MODELING

The first part, dedicated to the wake model, showed that dynamic wake models do describe the power generated by a wind farm more accurately than their steady-state counterparts under time-varying conditions (Chapter 2 to 4). The considered time-varying conditions are wind direction changes, turbine state changes, and turbulent wind speeds. This was tested with up to ten turbines in cases inspired by field measurements (Chapter 4 and 7). Centerline models have a significant computational advantage over multi-chain models while capturing the wake in a more true-to-model manner (Chapter 3). This comes at the cost of the ability to advect the different wake regions at different speeds. Another advantage of centerline models is that they can be coupled to a generic steady-state wake model, as introduced in Chapter 3 and further demonstrated in Chapter 4. This allows them to take advantage of further developments in the field of steady-state wake models. The developed models have proven to be capable to predict dynamic wake behavior at a low computational cost (Chapter 2 to 7).

8.1.2. STATE ESTIMATION

The second part of the thesis was dedicated to estimating the state of the derived surrogate model using measurements. To this end, an effort was made to utilize only widely available measurement signals, such as the power generated by the turbines and wind direction measurements at the turbine location. Chapter 5 portrayed how the flow field state of a particle-based engineering model can be estimated using an ensemble approach. The designed Ensemble Kalman Filter is able to estimate heterogeneous and time-varying flow field conditions and effectively couple the model with the true wind farm. A central element that makes this state estimation different from grid-based simulations is that particle-based simulation may have particles at different locations in each wake. The resulting state was then used to initialize the model for the control step

(Chapter 7). The work shows that measurements at the turbine locations are sufficient to estimate the dynamically changing flow field for wake steering control. Therefore, additional measurement equipment like LiDARs is not required to perform this kind of control. The framework further proved to add robustness towards sensor noise and biases (Chapter 7).

8.1.3. WIND FARM FLOW CONTROL

Wind farm control using dynamic surrogate models inherently leads to a time-dependent optimization problem. This has traditionally been addressed by calculating the energy generated over a prediction horizon. Depending on the degrees of freedom, this will lead to opportunistic behavior of the turbines, like the turnpike effect. Fewer degrees of freedom will lead to a conservative investment-return balancing behavior (Chapter 7). Steady-state models, however, neglect the transport delay and return the most aggressive control set points. Chapters 6 demonstrated a path between the two behaviors by synchronizing control actions with their effect and constructing several cost functions. This has proven to return steady-state-like performance under ideal conditions and better performance than steady-state controllers with disturbed wind direction measurements (Chapter 7). In a ten-turbine wind farm under turbulent and time-varying wind direction conditions, the proposed controller achieved an overall energy gain of 3 to 4 % over the baseline, depending on the controller settings. If disturbed and biased measurements are used, this performance decreases to 1.9 to 3 % over the greedy control baseline with the same measurements. The comparison to look-up table controllers shows that the closed-loop framework performance is more robust to disturbed measurements but can only match the performance in noise-free conditions.

8.1.4. OVERALL CONCLUSIONS

This thesis broke with the current practice to use steady-state engineering wake models for wind farm flow control by reincorporating wake dynamics. We showed that adding wake dynamics leads to a better prediction of turbine and wind farm power under time-varying conditions. However, given the same preview-less ideal-measurement conditions, we could not show that a dynamic wake description leads to a significant gain over the control set points derived by a steady-state model. This is mainly due to the difference in the nature of the cost function of maximizing power compared to energy. Maximizing energy entails recouping the investments of the control actions within a given time frame. This leads to more conservative control actions that are less sensitive to changes but also do not outperform the baseline by a large margin. Upon removing the intra-turbine delays from the cost function, the magnitude of the control set points becomes equivalent to the ones steady-state models predict. In a case study, we showed that the closed-loop design can be more robust to sensor noise and biases than a steady-state approach. This can be attributed to the state estimation.

We thereby conclude: Closed-loop wind farm flow control based on a dynamic engineering surrogate model leads to a more accurate and robust state estimation of the wind farm flow field but, given no preview, does not necessarily lead to a higher energy generation than what can be achieved with steady-state models.

8.2. RECOMMENDATIONS

Based on these conclusions, we formulate a set of recommendations for future work. First, Section 8.2.1 provides high-level recommendations on future wind farm flow control developments. This is followed by a set of recommendations individual to the three parts of this thesis: Wake modeling (Section 8.2.2), state estimation (Section 8.2.3), and control (Section 8.2.4).

8.2.1. OVERALL RECOMMENDATIONS

A limitation of this work was the use of readily available sensor data, which led to the limitation of using no preview during the optimization of the control decision. Future work should investigate the possibilities to eliminate this limitation as it can lead to significant gains in energy for both dynamic and steady-state models [95, 139, 152]. To this end, a mesoscale closed-loop wind farm flow control framework should be designed. This would need to dynamically describe and predict the changes in the atmosphere to derive the optimal control actions in a wind farm over the course of minutes to hours. A sensor fusion strategy is needed to connect information about wind farms, weather stations, and satellite data. This information must be collected in a dynamic model describing and predicting the flow conditions across the wind farm site and neighboring areas, e.g. [162]. Lastly, these flow conditions must be translated into the local wind farm flow field. This can be done using the heterogeneous flow field description of steady-state models [23], making them quasi-dynamic, or, a dynamic wake model, as presented in this thesis. The varying atmospheric conditions also require versatile wake models which can simulate turbine and farm wakes in stable, as well as unstable conditions. The proposed framework must further connect widely spread sensor data and bridge various space and time scales, from turbine-to-turbine to farm-to-farm. Recent work already explores some of these aspects, however, not for active control [163, 164].

8.2.2. DYNAMIC WAKE MODELING

The dynamic engineering model in this thesis is used in cases with up to ten turbines. Modern wind farms, however, consist of tens to hundreds of turbines, which requires additional modeling improvements. This includes aspects like the growing impact of turbulence throughout the farm [165], farm blockage [166], and gravity waves [167]. Steady-state models have and still face the same issues [86]. These corrections may also be applicable to dynamic models. Still, they may also face their own challenges, e.g., as the farm blockage depends on the turbine state and may, therefore, have its own dynamics in response to control set point changes [168]. This also requires a more thorough discussion of the relevant dynamics and time scales for wind farm flow control and integration of well-informed uncertainty bands covering unmodeled flow field effects. During the wake model development, the multi-OP-chain approach used in [1, 19] has been abandoned to lower the computational cost and improve the flexibility of the wake model. It may make sense to reintroduce the approach in the context of a stable Atmospheric boundary layer, where different parts of the wake experience different wind directions. The centerline approach further has the drawback that the wake advection speed is constant throughout all wake areas. This may be

addressed by changing how the extrapolation is done from the centerline to the surrounding areas.

The discussed changes mainly concern the wake model; however, there is also a need to improve the turbine model. This mainly concerns the estimation of loads experienced by the turbine. Recent work has proposed data-driven approaches to this problem [160, 161], which has not yet been integrated into the dynamic model. It may also be advantageous to find analytical relations between structural loads and easily obtainable quantities like yaw travel.

Dynamic models can furthermore be valuable for floating turbine and farm applications. First contributions have been made in this field with the goal of repositioning turbines to maximize the collective wind farm energy [15, 169]. Here, a dynamic model is necessary to model the impact of the changing turbine state onto the turbines position. Steady-state models have already been adapted to model floating turbines, which has for instance been used to optimize the mooring line configuration to achieve an ideal passive layout of the wind farm [170]. Dynamic wake models can build upon this by simulating the wind farm in time-varying conditions. dynamic models can further help to explore more futuristic approaches like floating islands of renewable energy [171].

Lastly, the computational performance of the code should be improved. This especially concerns the online application of the model for optimization purposes. This may entail rewriting the code to better use parallelization and the available hardware. In this context, it should be explored if the use of GPUs makes sense, as the model will feature a large number of particles that are advanced in parallel.

8.2.3. STATE AND PARAMETER ESTIMATION

The state estimation used in a closed-loop controller should further focus on available sensors and data sources. This entails deriving better models of sensor individual issues, such as noise and biases, based on their position within the wind farm. This can further be acknowledged in the correction, where, in the current design, all turbines are assumed to exhibit the same sensor noise and bias. Looping this information, obtained from the wake model, back may lead to an overall improvement of the state estimation. Another aspect is that there are three types of state estimation for engineering models: (i) flow field state estimation like the wind speed [83], (ii) wake model parameter estimation [51, 110], and (iii) wake location [90, 91]. These approaches largely use the same or similar available turbine measurements, e.g. the turbine power. Future work will have to combine these and will need to find a way to attribute prediction-measurement mismatches to the incorrectly modeled component. A decrease in power generated, for instance, may be caused by a wind speed reduction or a passing wake. This may be achieved by taking the varying observability of the components into account [172].

8.2.4. CONTROL AND OPTIMIZATION

Maximizing energy is a difficult cost function, as it aims for a supremum. This leads to an opportunistic turbine behavior for many degrees of freedom and conservative solutions for few. Future work will have to further explore how to balance the formulation of the cost function alongside the invested actuation to achieve realistic turbine actuation signals. This also entails taking uncertainty into

account. These may stem from the state estimation and historical data. It remains to be determined whether dynamic engineering wake models can significantly improve control decisions over steady-state or graph-based models to maximize the power generated by a farm. It is further desirable to find a manner to compare the performance of controllers outside of the study of dedicated cases to find more insight into their general behavior.

BIBLIOGRAPHY

- [1] M. Becker, B. Ritter, B. Doekemeijer, D. van der Hoek, U. Konigorski, D. Allaerts, and J.-W. van Wingerden. "The Revised FLORIDyn Model: Implementation of Heterogeneous Flow and the Gaussian Wake". In: *Wind Energy Science* 7.6 (Nov. 2022), pp. 2163–2179. issn: 2366-7451. doi: 10.5194/wes-7-2163-2022.
- [2] N. O. Jensen. *A Note on Wind Generator Interaction*. Roskilde, Denmark: Risø National Laboratory, 1983. isbn: 978-87-550-0971-4.
- [3] 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: *2014 American Control Conference*. June 2014, pp. 3128–3134. doi: 10.1109/ACC.2014.6859118.
- [4] P. Fleming, J. 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". In: *Wind Energy Science* 2.1 (May 2017), pp. 229–239. issn: 2366-7451. doi: 10.5194/wes-2-229-2017.
- [5] A. C. Kheirabadi and R. Nagamune. "A Quantitative Review of Wind Farm Control with the Objective of Wind Farm Power Maximization". In: *Journal of Wind Engineering and Industrial Aerodynamics* 192 (Sept. 2019), pp. 45–73. issn: 01676105. doi: 10.1016/j.jweia.2019.06.015.
- [6] L. E. Andersson, O. Anaya-Lara, J. O. Tande, K. O. Merz, and L. Imsland. "Wind Farm Control - Part I: A Review on Control System Concepts and Structures". In: *IET Renewable Power Generation* 15.10 (2021), pp. 2085–2108. issn: 1752-1424. doi: 10.1049/rpg2.12160.
- [7] J. D. Grunnet, M. Soltani, T. Knudsen, M. N. Kragelund, and T. Bak. "Aeolus Toolbox for Dynamics Wind Farm Model, Simulation and Control". In: *European Wind Energy Conference and Exhibition, EWEC 2010*. 2010.
- [8] S. Poushpas and W. Leithead. "Wind Farm Control through Dynamic Co-ordination of Wind Turbines Reference Power". In: Lisbon, Portugal, Nov. 2014. doi: 10.1201/b18973-101.
- [9] S. Frandsen, R. Barthelmie, S. Pryor, O. Rathmann, S. Larsen, J. Højstrup, and M. Thøgersen. "Analytical Modelling of Wind Speed Deficit in Large Offshore Wind Farms". In: *Wind Energy* 9.1-2 (Jan. 2006), pp. 39–53. issn: 1095-4244, 1099-1824. doi: 10.1002/we.189.
- [10] E. Bossanyi. "Combining Induction Control and Wake Steering for Wind Farm Energy and Fatigue Loads Optimisation". In: *IOP Publishing* 1037 (June 2018), p. 032011. issn: 1742-6596. doi: 10.1088/1742-6596/1037/3/032011.

- [11] Á. Jiménez, A. Crespo, and E. Migoya. "Application of a LES Technique to Characterize the Wake Deflection of a Wind Turbine in Yaw". In: *Wind Energy* 13.6 (2010), pp. 559–572. issn: 1099-1824. doi: 10.1002/we.380.
- [12] J. Ainslie. "Calculating the Flowfield in the Wake of Wind Turbines". In: *Journal of Wind Engineering and Industrial Aerodynamics* 27.1-3 (Jan. 1988), pp. 213–224. issn: 01676105. doi: 10.1016/0167-6105(88)90037-2.
- [13] 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". In: *Wind Energy* (Mar. 2017). doi: 10.1002/we.2093. (Visited on 11/16/2021).
- [14] C. R. Shapiro, D. F. Gayme, and C. Meneveau. "Modelling Yawed Wind Turbine Wakes: A Lifting Line Approach". In: *Journal of Fluid Mechanics* 841 (Apr. 2018), R1. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2018.75.
- [15] A. C. Kheirabadi and R. Nagamune. "A Low-Fidelity Dynamic Wind Farm Model for Simulating Time-Varying Wind Conditions and Floating Platform Motion". In: *Ocean Engineering* 234 (Aug. 2021), p. 109313. issn: 0029-8018. doi: 10.1016/j.oceaneng.2021.109313.
- [16] M. Bastankhah and F. Porté-Agel. "Experimental and Theoretical Study of Wind Turbine Wakes in Yawed Conditions". In: *Journal of Fluid Mechanics* 806 (Nov. 2016), pp. 506–541. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2016.595.
- [17] G. C. Larsen, ed. *Dynamic Wake Meandering Modeling*. Risø R, Report 1607. Roskilde: Risø National Laboratory, 2007. isbn: 978-87-550-3602-4.
- [18] H. A. Madsen, G. C. Larsen, T. J. Larsen, N. Troidborg, and R. Mikkelsen. "Calibration and Validation of the Dynamic Wake Meandering Model for Implementation in an Aeroelastic Code". In: *Journal of Solar Energy Engineering* 132.4 (Nov. 2010), p. 041014. issn: 0199-6231, 1528-8986. doi: 10.1115/1.4002555.
- [19] P. M. O. Gebraad and J. W. van Wingerden. "A Control-Oriented Dynamic Model for Wakes in Wind Plants". In: *Journal of Physics: Conference Series* 524 (June 2014), p. 012186. issn: 1742-6596. doi: 10.1088/1742-6596/524/1/012186.
- [20] 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". In: *Journal of Turbulence* 13 (Jan. 2012), N14. issn: 1468-5248. doi: 10.1080/14685248.2012.668191.
- [21] 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: *2015 American Control Conference (ACC)*. Chicago, Illinois, July 2015, pp. 1702–1708. doi: 10.1109/ACC.2015.7170978.

- [22] M. Becker. *Gaussian FLORIDyn, Matlab Implementation Belonging to the Paper: The Revised FLORIDyn Model: Implementation of Heterogeneous Flow and the Gaussian Wake*. 4TU.ResearchData. June 2022. doi: 10.4121/19867846.
- [23] A. Farrell, J. King, C. Draxl, R. Mudafort, N. Hamilton, C. J. Bay, P. Fleming, and E. Simley. "Design and Analysis of a Wake Model for Spatially Heterogeneous Flow". In: *Wind Energy Science* 6.3 (May 2021), pp. 737–758. issn: 2366-7451. doi: 10.5194/wes-6-737-2021.
- [24] A. Crespo and J. Hernández. "Turbulence Characteristics in Wind-Turbine Wakes". In: *Journal of Wind Engineering and Industrial Aerodynamics* 61.1 (June 1996), pp. 71–85. issn: 01676105. doi: 10.1016/0167-6105(95)00033-X.
- [25] D. Medici. "Experimental Studies of Wind Turbine Wakes - Power Optimisation and Meandering". PhD thesis. KTH Stockholm, 2005.
- [26] J. H. W. Lee and V. H. Chu. "Turbulent Round Jet in Coflow". In: *Turbulent Jets and Plumes*. Boston, MA: Springer US, 2003, pp. 179–209. isbn: 978-1-4613-5061-3 978-1-4615-0407-8. doi: 10.1007/978-1-4615-0407-8_6.
- [27] H. Vogel. "A Better Way to Construct the Sunflower Head". In: *Mathematical Biosciences* (1979). doi: 10.1016/0025-5564(79)90080-4.
- [28] G. Voronoi. "Nouvelles Applications Des Paramètres Continus à La Théorie Des Formes Quadratiques. Premier Mémoire. Sur Quelques Propriétés Des Formes Quadratiques Positives Parfaites." In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 1908.133 (Jan. 1908), pp. 97–102. issn: 0075-4102, 1435-5345. doi: 10.1515/crll.1908.133.97. (Visited on 07/21/2021).
- [29] G. Voronoi. "Nouvelles Applications Des Paramètres Continus à La Théorie Des Formes Quadratiques. Deuxième Mémoire. Recherches Sur Les Paralléloèdres Primitifs." In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 1908.134 (July 1908), pp. 198–287. issn: 0075-4102, 1435-5345. doi: 10.1515/crll.1908.134.198. (Visited on 07/21/2021).
- [30] G. I. Taylor. "The Spectrum of Turbulence". In: *Proceedings of the Royal Society of London* (1938). doi: 10.1098/rspa.1938.0032.
- [31] D. Schlipf, D. Trabucchi, O. Bischoff, M. Hofsäß, J. Mann, T. Mikkelsen, A. Rettenmeier, J. J. Trujillo, and M. Kühn. "Testing of Frozen Turbulence Hypothesis for Wind Turbine Applications with a Scanning LIDAR System". In: *Detaled Program*. June 2010, p. 5.
- [32] S. J. Andersen, J. N. Sørensen, and R. F. Mikkelsen. "Turbulence and Entrainment Length Scales in Large Wind Farms". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (Apr. 2017), p. 20160107. issn: 1364-503X, 1471-2962. doi: 10.1098/rsta.2016.0107.
- [33] C. Bak, F. Zahle, R. Bitsche, T. Kim, A. Yde, L. C. Henriksen, M. H. Hansen, J. P. A. A. Blasques, M. Gaunaa, and A. Natarajan. "The DTU 10-MW Reference Wind Turbine". In: May 2013.

- [34] J. Annoni, P. M. O. Gebraad, A. K. Scholbrock, P. A. Fleming, and J.-W. V. Wingerden. "Analysis of Axial-induction-based Wind Plant Control Using an Engineering and a High-order Wind Plant Model". In: *Wind Energy* 19.6 (June 2016), pp. 1135–1150. issn: 1095-4244, 1099-1824. doi: 10.1002/we.1891.
- [35] F. D. Bianchi, H. D. Battista, and R. J. Mantz. *Wind Turbine Control Systems: Principles, Modelling and Gain Scheduling Design*. 2007. isbn: 978-1-84628-492-2.
- [36] H. A. Madsen, T. J. Larsen, G. R. Pirrung, A. Li, and F. Zahle. "Implementation of the Blade Element Momentum Model on a Polar Grid and Its Aeroelastic Load Impact". In: *Wind Energy Science* 5.1 (Jan. 2020), pp. 1–27. issn: 2366-7451. doi: 10.5194/wes-5-1-2020.
- [37] J. Liew, A. M. Urbán, and S. J. Andersen. "Analytical Model for the Power--Yaw Sensitivity of Wind Turbines Operating in Full Wake". In: *Wind Energy Science* 5.1 (Mar. 2020), pp. 427–437. issn: 2366-7451. doi: 10.5194/wes-5-427-2020.
- [38] M. F. Howland, C. M. González, J. J. P. Martínez, J. B. Quesada, F. P. Larrañaga, N. K. Yadav, J. S. Chawla, and J. O. Dabiri. "Influence of Atmospheric Conditions on the Power Production of Utility-Scale Wind Turbines in Yaw Misalignment". In: *Journal of Renewable and Sustainable Energy* 12.6 (Nov. 2020), p. 063307. issn: 1941-7012. doi: 10.1063/5.0023746.
- [39] S. Emeis. *Wind Energy Meteorology: Atmospheric Physics for Wind Power Generation*. Second edition. Springer International Publishing, 2018. isbn: 978-3-319-72859-9.
- [40] A. Niayifar and F. Porté-Agel. "A New Analytical Model for Wind Farm Power Prediction". In: *Journal of Physics: Conference Series* 625 (June 2015), p. 012039. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/625/1/012039.
- [41] B. Doekemeijer, R. Storm, J. Schreiber, and Daanvanderhoek. *TU Delft-DataDrivenControl/FLORISSE_M: Stable Version from 2018-2019*. Zenodo. Jan. 2021. doi: 10.5281/ZENODO.4458669.
- [42] J. N. Sorensen and W. Z. Shen. "Numerical Modeling of Wind Turbine Wakes". In: *Journal of Fluids Engineering* 124.2 (May 2002), pp. 393–399. issn: 0098-2202. doi: 10.1115/1.1471361.
- [43] R. Ortega, F. Mancilla-David, and F. Jaramillo. "A Globally Convergent Wind Speed Estimator for Wind Turbine Systems". In: *International Journal of Adaptive Control and Signal Processing* 27.5 (2013), pp. 413–425. issn: 1099-1115. doi: 10.1002/acs.2319.
- [44] M. Becker. *SOWFA Simulation Setup Belonging to the Paper: The Revised FLORIDyn Model: Implementation of Heterogeneous Flow and the Gaussian Wake*. June 2022. doi: 10.4121/20026406.
- [45] M. J. van den Broek and J.-W. van Wingerden. "Dynamic Flow Modelling for Model-Predictive Wind Farm Control". In: *Journal of Physics: Conference Series*. Vol. 1618. 2020, p. 022023. doi: 10.1088/1742-6596/1618/2/022023.

- [46] M. Bastankhah, B. L. Welch, L. A. Martínez-Tossas, J. King, and P. Fleming. "Analytical Solution for the Cumulative Wake of Wind Turbines in Wind Farms". In: *Journal of Fluid Mechanics* 911 (Mar. 2021), A53. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2020.1037.
- [47] M. Becker, D. Allaerts, and J. W. van Wingerden. "FLORIDyn - A Dynamic and Flexible Framework for Real-Time Wind Farm Control". In: *Journal of Physics: Conference Series* 2265.3 (May 2022), p. 032103. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/3/032103.
- [48] T. Knudsen, T. Bak, and M. Svenstrup. "Survey of Wind Farm Control-Power and Fatigue Optimization: Survey of Wind Farm Control". In: *Wind Energy* 18.8 (Aug. 2015), pp. 1333-1351. issn: 10954244. doi: 10.1002/we.1760.
- [49] J. W. van Wingerden, P. A. Fleming, T. Göçmen, I. Eguinoa, B. M. Doekemeijer, K. Dykes, M. Lawson, E. Simley, J. King, D. Astrain, M. Iribas, C. L. Bottasso, J. Meyers, S. Raach, K. Kölle, and G. Giebel. "Expert Elicitation on Wind Farm Control". In: *Journal of Physics: Conference Series* 1618 (Sept. 2020), p. 022025. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/1618/2/022025.
- [50] 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". In: *Wind Energy* 19.1 (2016), pp. 95-114. doi: 10.1002/we.1822.
- [51] B. M. Doekemeijer, D. van der Hoek, and J. van Wingerden. "Closed-Loop Model-Based Wind Farm Control Using FLORIS under Time-Varying Inflow Conditions". In: *Renewable Energy* 156 (Aug. 2020), pp. 719-730. issn: 09601481. doi: 10.1016/j.renene.2020.04.007.
- [52] Y. Liu, A. K. Pamososuryo, R. M. G. Ferrari, and J. W. van Wingerden. "The Immersion and Invariance Wind Speed Estimator Revisited and New Results". In: *IEEE Control Systems Letters* 6 (2022), pp. 361-366. issn: 2475-1456. doi: 10.1109/LCSYS.2021.3076040. (Visited on 01/13/2022).
- [53] L. Masset, O. Bruls, and G. Kerschen. "Partition of the Circle in Cells of Equal Area and Shape". In: *Open Repository and Bibliography* (May 2011), p. 6.
- [54] V. V. Dighe, M. Becker, T. Göçmen, B. Sanderse, and J.-W. van Wingerden. "Sensitivity Analysis and Bayesian Calibration of a Dynamic Wind Farm Control Model: FLORIDyn". In: *Journal of Physics: Conference Series* 2265.2 (May 2022), p. 022062. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/2/022062.
- [55] M. Becker, M. Lejeune, P. Chatelain, D. Allaerts, R. Mudafort, and J.-W. van Wingerden. "A Dynamic Open-Source Model to Investigate Wake Dynamics in Response to Wind Farm Flow Control Strategies". In: *Wind Energy Science* 10.6 (2025). doi: 10.5194/wes-10-1055-2025.
- [56] G. Costanzo and G. Brindley. *Wind Energy in Europe - 2023 Statistics and the Outlook for 2024-2023*. Tech. rep. Wind Europe, Feb. 2024. (Visited on 10/18/2024).

- [57] 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". In: *Wind Energy* 23.8 (Aug. 2020), pp. 1739–1751. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2513.
- [58] M. Coquelet, L. Bricteux, M. Moens, and P. Chatelain. "A Reinforcement-learning Approach for Individual Pitch Control". In: *Wind Energy* 25.8 (Aug. 2022), pp. 1343–1362. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2734.
- [59] W. Munters and J. Meyers. "An Optimal Control Framework for Dynamic Induction Control of Wind Farms and Their Interaction with the Atmospheric Boundary Layer". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (Apr. 2017), p. 20160100. issn: 1364-503X, 1471-2962. doi: 10.1098/rsta.2016.0100.
- [60] P. Fleming, J. King, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. Van Dam, C. Bay, R. Mudafort, D. Jager, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake. "Continued Results from a Field Campaign of Wake Steering Applied at a Commercial Wind Farm -- Part 2". In: *Wind Energy Science* 5.3 (July 2020), pp. 945–958. issn: 2366-7451. doi: 10.5194/wes-5-945-2020.
- [61] B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann, G. Potenza, F. Calabretta, F. Fioretti, and J.-W. van Wingerden. "Field Experiment for Open-Loop Yaw-Based Wake Steering at a Commercial Onshore Wind Farm in Italy". In: *Wind Energy Science* 6.1 (Jan. 2021), pp. 159–176. issn: 2366-7451. doi: 10.5194/wes-6-159-2021.
- [62] P. Hulsman, M. Howland, T. Göçmen, and V. Petrović. "Assessing Closed-Loop Data-Driven Wind Farm Control Strategies within a Wind Tunnel". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032049. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032049.
- [63] P. Chatelain, S. Backaert, G. Winckelmans, and S. Kern. "Large Eddy Simulation of Wind Turbine Wakes". In: *Flow, Turbulence and Combustion* 91.3 (Oct. 2013), pp. 587–605. issn: 1386-6184, 1573-1987. doi: 10.1007/s10494-013-9474-8.
- [64] M. Moens, M. Lejeune, and P. Chatelain. "An Advanced Farm Flow Estimator for the Real-Time Evaluation of the Potential Wind Power of a down-Regulated Wind Farm". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032044. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032044.
- [65] J. Quick, R. N. King, G. Barter, and P. E. Hamlington. "Multifidelity Multiobjective Optimization for Wake-Steering Strategies". In: *Wind Energy Science* 7.5 (Sept. 2022), pp. 1941–1955. issn: 2366-7451. doi: 10.5194/wes-7-1941-2022.

- [66] M. J. van den Broek, B. Sanderse, and J.-W. van Wingerden. "Flow Modelling for Wind Farm Control: 2D vs. 3D". In: *Journal of Physics: Conference Series* 2265.3 (May 2022), p. 032086. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/3/032086.
- [67] Y. Marichal, I. De Visscher, P. Chatelain, and G. Winckelmans. "Towards Physics-Based Operational Modeling of the Unsteady Wind Turbine Response to Atmospheric and Wake-Induced Turbulence". In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012030. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012030.
- [68] D. Marten. *QBlade: A Modern Tool for the Aeroelastic Simulation of Wind Turbines*. 2020. doi: 10.14279/DEPOSITONCE-10646.
- [69] M. J. van den Broek, D. De Tavernier, P. Hulsman, D. van der Hoek, B. Sanderse, and J.-W. van Wingerden. "Free-Vortex Models for Wind Turbine Wakes under Yaw Misalignment -- a Validation Study on Far-Wake Effects". In: *Wind Energy Science* 8.12 (Dec. 2023), pp. 1909-1925. issn: 2366-7451. doi: 10.5194/wes-8-1909-2023.
- [70] T. Larsen, G. Larsen, M. Pedersen, K. Enevoldsen, and H. Madsen. "Validation of the Dynamic Wake Meander Model with Focus on Tower Loads". In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012027. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012027.
- [71] J. Jonkman, P. Doubrawa, N. Hamilton, J. Annoni, and P. Fleming. "Validation of FAST.Farm Against Large-Eddy Simulations". In: *Journal of Physics: Conference Series* 1037 (June 2018), p. 062005. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/1037/6/062005.
- [72] J. M. Jonkman, J. Annoni, G. Hayman, B. Jonkman, and A. Purkayastha. "Development of FAST.Farm: A New Multi-Physics Engineering Tool for Wind-Farm Design and Analysis". In: *35th Wind Energy Symposium*. Grapevine, Texas: American Institute of Aeronautics and Astronautics, Jan. 2017. isbn: 978-1-62410-456-5. doi: 10.2514/6.2017-0454.
- [73] J. Liew, T. Göçmen, A. W. H. Lio, and G. C. Larsen. "Extending the Dynamic Wake Meandering Model in HAWC2Farm: A Comparison with Field Measurements at the Lillgrund Wind Farm". In: *Wind Energy Science* 8.9 (Sept. 2023), pp. 1387-1402. issn: 2366-7451. doi: 10.5194/wes-8-1387-2023.
- [74] R. Braunbehrens, J. Schreiber, and C. L. Bottasso. "Application of an Open-Loop Dynamic Wake Model with High-Frequency SCADA Data". In: *Journal of Physics: Conference Series* 2265.2 (May 2022), p. 022031. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/2/022031.
- [75] M. Lejeune, M. Moens, and P. Chatelain. "A Meandering-Capturing Wake Model Coupled to Rotor-Based Flow-Sensing for Operational Wind Farm Flow Prediction". In: *Frontiers in Energy Research* 10 (July 2022), p. 884068. issn: 2296-598X. doi: 10.3389/fenrg.2022.884068.

- [76] B. Foloppe, W. Munters, S. Buckingham, L. Vandeveld, and J. van Beeck. "Development of a Dynamic Wake Model Accounting for Wake Advection Delays and Mesoscale Wind Transients". In: *Journal of Physics: Conference Series* 2265.2 (May 2022), p. 022055. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/2/022055.
- [77] V. Kipke and C. Sourkounis. "Three-Dimensional Dynamic Wake Model for Real-Time Wind Farm Simulation". In: *2024 32nd Mediterranean Conference on Control and Automation (MED)*. Chania - Crete, Greece: IEEE, June 2024, pp. 808–815. isbn: 9798350395440. doi: 10.1109/MED61351.2024.10566140.
- [78] NREL. "FLORIS. Version 3.4". In: *GitHub repository* (Nov. 2023).
- [79] M. M. Pedersen, A. Meyer Forsting, P. van der Laan, R. Riva, L. A. Alcayaga Román, J. Criado Risco, M. Friis-Møller, J. Quick, J. P. Schøler Christiansen, R. Valotta Rodrigues, B. T. Olsen, and P.-E. Réthoré. "PyWake 2.5.0: An Open-Source Wind Farm Simulation Tool". In: *DTU Wind, Technical University of Denmark* (Feb. 2023).
- [80] J. Schmidt, L. Vollmer, M. Dörenkämper, and B. Stoevesandt. "FOXES: Farm Optimization and eXtended Yield Evaluation Software". In: *Journal of Open Source Software* 8.86 (June 2023), p. 5464. issn: 2475-9066. doi: 10.21105/joss.05464.
- [81] A. Sterle, C. A. Hans, and J. Raisch. "Model Predictive Control of Wakes for Wind Farm Power Tracking". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032005. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032005.
- [82] Y. Miao, M. N. Soltani, A. Hajizadeh, and S. Simani. "Artificial Neural Network-based Wake Steering Control under the Time-varying Inflow". In: *2024 10th International Conference on Control, Decision and Information Technologies (CoDIT)*. Vallette, Malta: IEEE, July 2024, pp. 1988–1993. isbn: 9798350373974. doi: 10.1109/CoDIT62066.2024.10708147.
- [83] M. Becker, D. Allaerts, and J.-W. van Wingerden. "Ensemble-Based Flow Field Estimation Using the Dynamic Wind Farm Model FLORIDyn". In: *Energies* 15.22 (Nov. 2022), p. 8589. issn: 1996-1073. doi: 10.3390/en15228589.
- [84] U. Ciri, M. A. Rotea, and S. Leonardi. "Model-Free Control of Wind Farms: A Comparative Study between Individual and Coordinated Extremum Seeking". In: *Renewable Energy* 113 (Dec. 2017), pp. 1033–1045. issn: 09601481. doi: 10.1016/j.renene.2017.06.065.
- [85] H. Zong and F. Porté-Agel. "A Momentum-Conserving Wake Superposition Method for Wind Farm Power Prediction". In: *Journal of Fluid Mechanics* 889 (Apr. 2020), A8. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2020.77.
- [86] C. J. Bay, P. Fleming, B. Doekemeijer, J. King, M. Churchfield, and R. Mudafort. "Addressing Deep Array Effects and Impacts to Wake Steering with the Cumulative-Curl Wake Model". In: *Wind Energy Science* 8.3 (Mar. 2023), pp. 401–419. issn: 2366-7451. doi: 10.5194/wes-8-401-2023.

- [87] W. J. M. Rankine. "On the Mechanical Principles of the Action of Propellers". In: *Transactions of the Institution of Naval Architects* 6 (1865), p. 13.
- [88] R. E. Froude. "On the Part Played in Propulsion by Differences of Fluid Pressure". In: *Transactions of the Institution of Naval Architects* 30 (1889), p. 390.
- [89] S. Tamaro, F. Campagnolo, and C. L. Bottasso. "On the Power and Control of a Misaligned Rotor – beyond the Cosine Law". In: *Wind Energy Science* 9.7 (July 23, 2024), pp. 1547–1575. issn: 2366-7451. doi: 10.5194/wes-9-1547-2024. url: <https://wes.copernicus.org/articles/9/1547/2024/>.
- [90] R. Braunbehrens, S. Tamaro, and C. L. Bottasso. "Towards the Multi-Scale Kalman Filtering of Dynamic Wake Models: Observing Turbulent Fluctuations and Wake Meandering". In: *Journal of Physics: Conference Series* 2505.1 (May 2023), p. 012044. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2505/1/012044.
- [91] J. Di Cave, R. Braunbehrens, J. Krause, A. Guilloré, and C. L. Bottasso. "Closed-Loop Coupling of a Dynamic Wake Model with a Wind Inflow Estimator". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032034. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032034.
- [92] S. Kanev. "Dynamic Wake Steering and Its Impact on Wind Farm Power Production and Yaw Actuator Duty". In: *Renewable Energy* 146 (Feb. 2020), pp. 9–15. issn: 09601481. doi: 10.1016/j.renene.2019.06.122.
- [93] P. A. Fleming, A. P. J. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doeke-meijer, and R. Madafort. "Serial-Refine Method for Fast Wake-Steering Yaw Optimization". In: *Journal of Physics: Conference Series* 2265.3 (May 2022), p. 032109. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2265/3/032109.
- [94] S. Knoop. *Wind - Lidar Wind Profiles Measured at North Sea Wind Farm TenneT Platforms 1 Second Raw Data*. Nov. 2019.
- [95] M. J. van den Broek, M. Becker, B. Sanderse, and J.-W. van Wingerden. "Dynamic Wind Farm Flow Control Using Free-Vortex Wake Models". In: *Wind Energy Science* 9.3 (Mar. 2024), pp. 721–740. issn: 2366-7451. doi: 10.5194/wes-9-721-2024.
- [96] J. Annoni, C. Bay, K. Johnson, E. Dall'Anese, E. Quon, T. Kemper, and P. Fleming. "Wind Direction Estimation Using SCADA Data with Consensus-Based Optimization". In: *Wind Energy Science* 4.2 (June 2019), pp. 355–368. issn: 2366-7451. doi: 10.5194/wes-4-355-2019.
- [97] D. van der Hoek, M. Sinner, E. Simley, L. Pao, and J.-W. van Wingerden. "Estimation of the Ambient Wind Field From Wind Turbine Measurements Using Gaussian Process Regression". In: *2021 American Control Conference (ACC)*. New Orleans, LA, USA: IEEE, May 2021, pp. 558–563. isbn: 978-1-66544-197-1. doi: 10.23919/ACC50511.2021.9483088.

- [98] M. F. Howland, H. M. Johlas, J. B. Quesada, J. J. Pena Martinez, W. Zhong, and F. P. Larranaga. "On the Impact of the Yaw Update Frequency and Wind Direction Forecasting on Open-Loop Wake Steering Control". In: *2022 American Control Conference (ACC)*. Atlanta, GA, USA: IEEE, June 2022, pp. 4218–4223. isbn: 978-1-66545-196-3. doi: 10.23919/ACC53348.2022.9867443.
- [99] M. Lejeune, A. Frère, M. Moens, and P. Chatelain. "Are Steady-State Wake Models and Lookup Tables Sufficient to Design Profitable Wake Steering Strategies? A Large Eddy Simulation Investigation". In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092075. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/9/092075.
- [100] L. Martinez, S. Leonardi, M. Churchfield, and P. Moriarty. "A Comparison of Actuator Disk and Actuator Line Wind Turbine Models and Best Practices for Their Use". In: *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. Nashville, Tennessee: American Institute of Aeronautics and Astronautics, Jan. 2012. isbn: 978-1-60086-936-5. doi: 10.2514/6.2012-900.
- [101] C. R. Shapiro, D. F. Gayme, and C. Meneveau. "Filtered Actuator Disks: Theory and Application to Wind Turbine Models in Large Eddy Simulation". In: *Wind Energy* 22.10 (Oct. 2019), pp. 1414–1420. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2376.
- [102] W. H. Lio, G. C. Larsen, and G. R. Thorsen. "Dynamic Wake Tracking Using a Cost-Effective LiDAR and Kalman Filtering: Design, Simulation and Full-Scale Validation". In: *Renewable Energy* 172 (July 2021), pp. 1073–1086. issn: 09601481. doi: 10.1016/j.renene.2021.03.081.
- [103] E. Simley, P. Fleming, and J. King. "Design and Analysis of a Wake Steering Controller with Wind Direction Variability". In: *Wind Energy Science* 5.2 (Apr. 2020), pp. 451–468. issn: 2366-7451. doi: 10.5194/wes-5-451-2020.
- [104] K. Heck, H. Johlas, and M. Howland. "Modelling the Induction, Thrust and Power of a Yaw-Misaligned Actuator Disk". In: *Journal of Fluid Mechanics* 959 (Mar. 2023), A9. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.129.
- [105] G. M. Starke, C. Meneveau, J. R. King, and D. F. Gayme. "A Dynamic Model of Wind Turbine Yaw for Active Farm Control". In: *Wind Energy* (Dec. 2023), we.2884. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2884.
- [106] M. Abkar, J. N. Sørensen, and F. Porté-Agel. "An Analytical Model for the Effect of Vertical Wind Veer on Wind Turbine Wakes". In: *Energies* 11.7 (July 13, 2018), p. 1838. issn: 1996-1073. doi: 10.3390/en11071838. url: <https://www.mdpi.com/1996-1073/11/7/1838>.
- [107] D. Bastine, B. Witha, M. Wächter, and J. Peinke. "Towards a Simplified DynamicWake Model Using POD Analysis". In: *Energies* 8.2 (Jan. 2015), pp. 895–920. issn: 1996-1073. doi: 10.3390/en8020895.

- [108] J. Gutknecht, M. Becker, C. Muscari, T. Lutz, and J.-W. van Wingerden. "Scaling DMD Modes for Modeling Dynamic Induction Control Wakes in Various Wind Speeds". In: *2023 IEEE Conference on Control Technology and Applications (CCTA)*. Bridgetown, Barbados: IEEE, Aug. 2023, pp. 574–580. isbn: 9798350335446. doi: 10.1109/CCTA54093.2023.10252400.
- [109] M. J. van den Broek, D. De Tavernier, B. Sanderse, and J.-W. van Wingerden. "Adjoint Optimisation for Wind Farm Flow Control with a Free-Vortex Wake Model". In: *Renewable Energy* (Nov. 2022), S0960148122016226. issn: 09601481. doi: 10.1016/j.renene.2022.10.120.
- [110] M. F. Howland, A. S. Ghatge, S. K. Lele, and J. O. Dabiri. "Optimal Closed-Loop Wake Steering -- Part 1: Conventionally Neutral Atmospheric Boundary Layer Conditions". In: *Wind Energy Science* 5.4 (Oct. 2020), pp. 1315–1338. issn: 2366-7443. doi: 10.5194/wes-5-1315-2020.
- [111] G. Evensen. "The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation". In: *Ocean Dynamics* 53.4 (Nov. 2003), pp. 343–367. issn: 1616-7341, 1616-7228. doi: 10.1007/s10236-003-0036-9. (Visited on 12/08/2020).
- [112] O.-P. Tossavainen, J. Percelay, A. Tinka, Q. Wu, and A. M. Bayen. "Ensemble Kalman Filter Based State Estimation in 2D Shallow Water Equations Using Lagrangian Sensing and State Augmentation". In: *2008 47th IEEE Conference on Decision and Control*. Cancun, Mexico: IEEE, 2008, pp. 1783–1790. isbn: 978-1-4244-3123-6. doi: 10.1109/CDC.2008.4738999.
- [113] J. Du, J. Zhu, F. Fang, C. C. Pain, and I. M. Navon. "Ensemble Data Assimilation Applied to an Adaptive Mesh Ocean Model: Ensemble Data Assimilation Applied to an Adaptive Mesh Ocean Model". In: *International Journal for Numerical Methods in Fluids* 82.12 (Dec. 2016), pp. 997–1009. issn: 02712091. doi: 10.1002/flid.4247.
- [114] B. M. Doekemeijer, S. Boersma, L. Y. Pao, and J. W. van Wingerden. "Ensemble Kalman Filtering for Wind Field Estimation in Wind Farms". In: *2017 American Control Conference (ACC)*. May 2017, pp. 19–24. doi: 10.23919/ACC.2017.7962924.
- [115] P. Kumar Jain, K. Mandli, I. Hoteit, O. Knio, and C. Dawson. "Dynamically Adaptive Data-Driven Simulation of Extreme Hydrological Flows". In: *Ocean Modelling* 122 (Feb. 2018), pp. 85–103. issn: 14635003. doi: 10.1016/j.ocemod.2017.12.004.
- [116] C. R. Shapiro, G. M. Starke, C. Meneveau, and D. F. Gayme. "A Wake Modeling Paradigm for Wind Farm Design and Control". In: *Energies* 12.15 (Aug. 2019), p. 2956. issn: 1996-1073. doi: 10.3390/en12152956.
- [117] C. Sampson, A. Carrassi, A. Aydoğdu, and C. K. Jones. "Ensemble Kalman Filter for Nonconservative Moving Mesh Solvers with a Joint Physics and Mesh Location Update". In: *Quarterly Journal of the Royal Meteorological Society* 147.736 (Apr. 2021), pp. 1539–1561. issn: 0035-9009, 1477-870X. doi: 10.1002/qj.3980.
- [118] S. J. Julier and J. K. Uhlmann. "New Extension of the Kalman Filter to Nonlinear Systems". In: *AeroSense '97*. Ed. by I. Kadar. Orlando, FL, USA, July 1997, p. 182. doi: 10.1117/12.280797. (Visited on 10/03/2022).

- [119] O. Talagrand and P. Courtier. "Variational Assimilation of Meteorological Observations With the Adjoint Vorticity Equation. I: Theory: VARIATIONAL ASSIMILATION. I: THEORY". In: *Quarterly Journal of the Royal Meteorological Society* 113.478 (Oct. 1987), pp. 1311–1328. issn: 00359009. doi: 10.1002/qj.49711347812.
- [120] E. Wan and R. Van Der Merwe. "The Unscented Kalman Filter for Nonlinear Estimation". In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*. Lake Louise, Alta., Canada: IEEE, 2000, pp. 153–158. isbn: 978-0-7803-5800-3. doi: 10.1109/ASSPCC.2000.882463.
- [121] P. Bauweraerts and J. Meyers. "Reconstruction of Turbulent Flow Fields from Lidar Measurements Using Large-Eddy Simulation". In: *Journal of Fluid Mechanics* 906 (Jan. 2021), A17. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2020.805.
- [122] S. Skachko, Q. Errera, R. Ménard, Y. Christophe, and S. Chabrilat. "Comparison of the Ensemble Kalman Filter and 4D-Var Assimilation Methods Using a Stratospheric Tracer Transport Model". In: *Geoscientific Model Development* 7.4 (July 2014), pp. 1451–1465. issn: 1991-9603. doi: 10.5194/gmd-7-1451-2014.
- [123] A. Betz. *Introduction to the Theory of Flow Machines*. Burlington: Elsevier Science, 2014. isbn: 978-1-4831-8090-8.
- [124] G. Evensen. "Sequential Data Assimilation with a Nonlinear Quasi-Geostrophic Model Using Monte Carlo Methods to Forecast Error Statistics". In: *Journal of Geophysical Research: Oceans* 99.C5 (1994), pp. 10143–10162. issn: 2156-2202. doi: 10.1029/94JC00572.
- [125] G. Burgers, P. J. van Leeuwen, and G. Evensen. "Analysis Scheme in the Ensemble Kalman Filter". In: *Monthly Weather Review* 126.6 (June 1998), pp. 1719–1724. issn: 1520-0493, 0027-0644. doi: 10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2.
- [126] R. Petrie. "Localization in the Ensemble Kalman Filter". In: 2008.
- [127] G. Gaspari and S. E. Cohn. "Construction of Correlation Functions in Two and Three Dimensions". In: *Quarterly Journal of the Royal Meteorological Society* 125.554 (Jan. 1999), pp. 723–757. issn: 00359009, 1477870X. doi: 10.1002/qj.49712555417.
- [128] A. C. Lorenc. "The Potential of the Ensemble Kalman Filter for NWP---a Comparison with 4D-Var". In: *Quarterly Journal of the Royal Meteorological Society* 129.595 (Oct. 2003), pp. 3183–3203. issn: 00359009, 1477870X. doi: 10.1256/qj.02.132.
- [129] National Renewable Energy Laboratory. "Simulator for Offshore Wind Farm Applications". In: *GitHub repository* (Nov. 2020).
- [130] M. Becker. *Dataset Belonging to the Paper: Ensemble Based Flow Field Estimation Using the Dynamic Wake Model FLORIDyn*. Sept. 2022.

- [131] M. Becker, D. Allaerts, and J.-W. van Wingerden. "Time-Shifted Cost Function Design for More Efficient Dynamic Wind Farm Flow Control*". In: *2024 IEEE Conference on Control Technology and Applications (CCTA)*. Newcastle upon Tyne, United Kingdom: IEEE, Aug. 2024, pp. 440–445. isbn: 9798350370942. doi: 10.1109/CCTA60707.2024.10666535.
- [132] J. Meyers, C. Bottasso, K. Dykes, P. Fleming, P. Gebraad, G. Giebel, T. Göçmen, and J.-W. van Wingerden. "Wind Farm Flow Control: Prospects and Challenges". In: *Wind Energy Science* 7.6 (Nov. 2022), pp. 2271–2306. issn: 2366-7451. doi: 10.5194/wes-7-2271-2022.
- [133] M. F. Howland, S. K. Lele, and J. O. Dabiri. "Wind Farm Power Optimization through Wake Steering". In: *Proceedings of the National Academy of Sciences* 116.29 (July 2019), pp. 14495–14500. issn: 0027-8424, 1091-6490. doi: 10.1073/pnas.1903680116.
- [134] S. Boersma, B. Doekemeijer, M. Vali, J. Meyers, and J.-W. van Wingerden. "A Control-Oriented Dynamic Wind Farm Model: WFSim". In: *Wind Energy Science* 3.1 (Mar. 2018), pp. 75–95. issn: 2366-7451. doi: 10.5194/wes-3-75-2018.
- [135] R. Dorfman, P. A. Samuelson, and R. M. Solow. *Linear Programming and Economic Analysis*. New York: Dover Publications, 1987. isbn: 978-0-486-65491-1.
- [136] S. Macrì, T. Duc, A. Leroy, N. Girard, and S. Aubrun. "Experimental Analysis of Time Delays in Wind Turbine Wake Interactions". In: *Journal of Physics: Conference Series* 1618.6 (Sept. 2020), p. 062058. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/1618/6/062058.
- [137] S. Emeis. *Wind Energy Meteorology*. Green Energy and Technology. Cham: Springer International Publishing, 2018. isbn: 978-3-319-72858-2 978-3-319-72859-9. doi: 10.1007/978-3-319-72859-9.
- [138] MathWorks. *MATLAB - Optimization Toolbox Version 9.5*. The MathWorks Inc. Natick, Massachusetts, 2023.
- [139] B. A. M. Sengers, A. Rott, E. Simley, M. Sinner, G. Steinfeld, and M. Kühn. "Increased Power Gains from Wake Steering Control Using Preview Wind Direction Information". In: *Wind Energy Science* 8.11 (Nov. 2023), pp. 1693–1710. issn: 2366-7451. doi: 10.5194/wes-8-1693-2023.
- [140] M. Becker, M. J. van den Broek, D. Allaerts, and J.-W. van Wingerden. *Closed-Loop Model-Predictive Wind Farm Flow Control under Time-Varying Inflow Using FLORIDyn*. Dec. 2024.
- [141] UNECE. *Carbon Neutrality in the UNECE Region: Integrated Life-cycle Assessment of Electricity Sources*. Tech. rep. Geneva: UNITED NATIONS ECONOMIC COMMISSION FOR EUROPE, 2022.
- [142] A. Rott, S. Boersma, J.-W. van Wingerden, and M. Kühn. "Dynamic Flow Model for Real-Time Application in Wind Farm Control". In: *Journal of Physics: Conference Series* 854 (May 2017), p. 012039. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/854/1/012039.
- [143] J. Liew. "Dynamic Modelling of Wind Farms for Closed-Loop Control". PhD thesis. Technical University of Denmark, 2022. doi: 10.11581/DTU.00000263.

- [144] T. Duc, O. Coupiac, N. Girard, G. Giebel, and T. Göçmen. "Local Turbulence Parameterization Improves the Jensen Wake Model and Its Implementation for Power Optimization of an Operating Wind Farm". In: *Wind Energy Science* 4.2 (May 2019), pp. 287–302. issn: 2366-7451. doi: 10.5194/wes-4-287-2019.
- [145] M. F. Howland, A. S. Ghaté, J. B. Quesada, J. J. Pena Martínez, W. Zhong, F. P. Larrañaga, S. K. Lele, and J. O. Dabiri. "Optimal Closed-Loop Wake Steering -- Part 2: Diurnal Cycle Atmospheric Boundary Layer Conditions". In: *Wind Energy Science* 7.1 (Feb. 2022), pp. 345–365. issn: 2366-7451. doi: 10.5194/wes-7-345-2022.
- [146] I. Sood and J. Meyers. "Development and Validation of a Large Eddy Simulation Based Virtual Environment for Optimal Wind Farm Control". PhD thesis. KU Leuven, May 2023.
- [147] P. Bachant, P. Ireland, B. Burrows, C. Qiao, J. Duncan, D. Zheng, and M. Dua. "Development and Validation of a Hybrid Data-Driven Model-Based Wake Steering Controller and Its Application at a Utility-Scale Wind Plant". In: *Wind Energy Science* 9.11 (Nov. 2024), pp. 2235–2259. issn: 2366-7451. doi: 10.5194/wes-9-2235-2024.
- [148] J. P. Goit and J. Meyers. "Optimal Control of Energy Extraction in Wind-Farm Boundary Layers". In: *Journal of Fluid Mechanics* 768 (Apr. 2015), pp. 5–50. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2015.70.
- [149] W. Munters and J. Meyers. "Dynamic Strategies for Yaw and Induction Control of Wind Farms Based on Large-Eddy Simulation and Optimization". In: *Energies* 11.1 (Jan. 2018), p. 177. doi: 10.3390/en11010177.
- [150] J. Goit, W. Munters, and J. Meyers. "Optimal Coordinated Control of Power Extraction in LES of a Wind Farm with Entrance Effects". In: *Energies* 9.1 (Jan. 2016), p. 29. issn: 1996-1073. doi: 10.3390/en9010029.
- [151] N. Janssens and J. Meyers. "Towards Real-Time Optimal Control of Wind Farms Using Large-Eddy Simulations". In: *Wind Energy Science* 9.1 (Jan. 2024), pp. 65–95. issn: 2366-7451. doi: 10.5194/wes-9-65-2024.
- [152] E. Simley, P. Fleming, J. King, and M. Sinner. "Wake Steering Wind Farm Control With Preview Wind Direction Information". In: *2021 American Control Conference (ACC)*. New Orleans, LA, USA: IEEE, May 2021, pp. 1783–1789. isbn: 978-1-66544-197-1. doi: 10.23919/ACC50511.2021.9483008.
- [153] J. Schreiber, C. L. Bottasso, and M. Bertelè. "Field Testing of a Local Wind Inflow Estimator and Wake Detector". In: *Wind Energy Science* 5.3 (July 2020), pp. 867–884. issn: 2366-7451. doi: 10.5194/wes-5-867-2020.
- [154] G. Evensen. *Data Assimilation: The Ensemble Kalman Filter*. 2. ed. Dordrecht: Springer, 2009. isbn: 978-3-642-03711-5 978-3-642-03710-8 978-3-642-42476-2.
- [155] M.-G. Kim and P. H. Dalhoff. "Yaw Systems for Wind Turbines -- Overview of Concepts, Current Challenges and Design Methods". In: *Journal of Physics: Conference Series* 524 (June 2014), p. 012086. issn: 1742-6596. doi: 10.1088/1742-6596/524/1/012086.

- [156] F. Bernardoni, U. Ciri, M. A. Rotea, and S. Leonardi. "Identification of Wind Turbine Clusters for Effective Real Time Yaw Control Optimization". In: *Journal of Renewable and Sustainable Energy* 13.4 (July 2021), p. 043301. issn: 1941-7012. doi: 10.1063/5.0036640.
- [157] IEEE Neural Networks Council, ed. *Proceedings / 1995 IEEE International Conference on Neural Networks, the University of Western Australia, Perth, Western Australia, 27 November - 1 December 1995*. Piscataway, NJ: IEEE Service Center, 1995. isbn: 978-0-7803-2768-9 978-0-7803-2769-6 978-0-7803-2770-2.
- [158] M. E. H. Pedersen. *Good Parameters for Particle Swarm Optimization*. Tech. rep. Technical Report no. HL1001. Luxembourg: Hvass Laboratories, 2010.
- [159] E. Mezura-Montes and C. A. Coello Coello. "Constraint-Handling in Nature-Inspired Numerical Optimization: Past, Present and Future". In: *Swarm and Evolutionary Computation* 1.4 (Dec. 2011), pp. 173–194. issn: 22106502. doi: 10.1016/j.swevo.2011.10.001.
- [160] J. Liew, R. Riva, M. Friis-Møller, and T. Göçmen. "Wind Farm Control Optimisation Under Load Constraints Via Surrogate Modelling". In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092039. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/9/092039.
- [161] A. Guilloiré, F. Campagnolo, and C. L. Bottasso. "A Control-Oriented Load Surrogate Model Based on Sector-Averaged Inflow Quantities: Capturing Damage for Unwaked, Waked, Wake-Steering and Curtailed Wind Turbines". In: *Journal of Physics: Conference Series* 2767.3 (June 2024), p. 032019. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/3/032019.
- [162] U. N. S. Foundation and National Center for Atmospheric Research. *Weather Research and Forecasting (WRF) Model*.
- [163] S. Porchetta, M. F. Howland, R. Borgers, S. Buckingham, and W. Munters. *Annual Wake Impacts in and between Wind Farm Clusters Modelled by a Mesoscale Numerical Weather Prediction Model and Fast-Running Engineering Models*. June 2024. doi: 10.5194/wes-2024-58.
- [164] T. Ivanova, S. Porchetta, S. Buckingham, J. Van Beeck, and W. Munters. *Improving Wind and Power Predictions via Four-Dimensional Data Assimilation in the WRF Model: Case Study of Storms in February 2022 at Belgian Offshore Wind Farms*. Mar. 2024. doi: 10.5194/wes-2023-177.
- [165] P. Argyle, S. Watson, C. Montavon, I. Jones, and M. Smith. "Modelling Turbulence Intensity within a Large Offshore Wind Farm". In: *Wind Energy* 21.12 (Dec. 2018), pp. 1329–1343. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2257.
- [166] J. Bleeg, M. Purcell, R. Ruisi, and E. Traiger. "Wind Farm Blockage and the Consequences of Neglecting Its Impact on Energy Production". In: *Energies* 11.6 (June 2018), p. 1609. issn: 1996-1073. doi: 10.3390/en11061609.
- [167] D. Allaerts and J. Meyers. "Boundary-Layer Development and Gravity Waves in Conventionally Neutral Wind Farms". In: *Journal of Fluid Mechanics* 814 (Mar. 2017), pp. 95–130. issn: 0022-1120, 1469-7645. doi: 10.1017/jfm.2017.11.

- [168] E. Bossanyi and J. Bleeg. "How Do Wind Farm Blockage and Axial Induction Control Interact?" In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092027. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/9/092027.
- [169] A. C. Kheirabadi and R. Nagamune. "Real-Time Relocation of Floating Off-shore Wind Turbines for Power Maximization Using Distributed Economic Model Predictive Control". In: *2021 American Control Conference (ACC)*. New Orleans, LA, USA: IEEE, May 2021, pp. 3077–3081. isbn: 978-1-66544-197-1. doi: 10.23919/ACC50511.2021.9483056.
- [170] Y. R. Alkarem, K. Huguenard, A. S. Verma, D. Van Binsbergen, E. Bachynski-Polić, and A. R. Nejad. "Passive Mooring-based Turbine Repositioning Technique for Wake Steering in Floating Offshore Wind Farms". In: *Journal of Physics: Conference Series* 2767.9 (June 2024), p. 092056. issn: 1742-6588, 1742-6596. doi: 10.1088/1742-6596/2767/9/092056.
- [171] P. J. Stuyfzand and J. W. Kappelhof. "Floating, High-Capacity Desalting Islands on Renewable Multi-Energy Supply". In: *Desalination* 177.1-3 (June 2005), pp. 259–266. issn: 00119164. doi: 10.1016/j.desal.2004.12.011.
- [172] B. Doekemeijer and J.-W. van Wingerden. "Observability of the Ambient Conditions in Model-based Estimation for Wind Farm Control: A Focus on Static Models". In: *Wind Energy* 23.9 (Sept. 2020), pp. 1777–1791. issn: 1095-4244, 1099-1824. doi: 10.1002/we.2495.



EPILOGUE



Mistakes and errors are inherent to the process of any work conducted, mine being no exception to it. By acknowledging this fact, mistakes can become more than annoyances and evolve into exciting lessons. Filipe Fortes captured the experience regarding software development as follows:



Debugging is like being the detective in a crime movie where you are also the murderer.

Filipe Fortes

In this epilogue, I present you some film stills of the “crime movie” that I have been a part of in recent years. Each figure represents a little riddle solved and a lesson learned after an investigation. To you, dear reader, best of luck with failing, and enjoy the process of learning from mistakes as much as possible.







CURRICULUM VITAE

Marcus Becker

12-01-1994 Born in Freiburg im Breisgau, Germany.

EDUCATION

2020-2024 Doctoral degree in Systems and Control
Delft University of Technology, Delft, Netherlands
Thesis: A comprehensive approach to closed-loop wind
farm flow control using FLORIDyn
Promotor: Prof. dr. ir. J.W. van Wingerden
Copromotor: Prof. dr. D. Allaerts

2018-2020 Master of Science in simulation and control of mechatronic systems
Technical University of Darmstadt, Darmstadt, Germany
Hosted by Delft University of Technology, Delft, Netherlands
Thesis: Maturing FLORIS towards a Dynamic Wind Farm
Model

2016 Erasmus+
Chalmers University of Technology, Göteborg, Sweden

2013-2017 Bachelor of Science in Mechatronics
Technical University of Darmstadt, Darmstadt, Germany

2004-2012 Gymnasium
Heimschule Lender, Sasbach, Germany

AWARDS

2023 Outstanding TA Award
2025 Best Wind & Ocean Energy Research Visualisation Award 2024

LIST OF PUBLICATIONS

JOURNAL PUBLICATIONS

6. *M. Becker*, M. J. van den Broek, D. Allaerts, and J.W. van Wingerden, *Closed-Loop Model-Predictive Wind Farm Flow Control under Time-Varying Inflow Using FLORIDyn*, accepted for publication in *Wind Energy* (2025).
5. *M. Becker*, M. Lejeune, P. Chatelain, D. Allaerts, R. Mudafort, and J.W. van Wingerden, *A Dynamic Open-Source Model to Investigate Wake Dynamics in Response to Wind Farm Flow Control Strategies*, *Wind Energy Science* (2025). doi: 10.5194/wes-10-1055-2025
4. M. J. van den Broek, *M. Becker*, B. Sanderse, and J.W. van Wingerden, *Dynamic Wind Farm Flow Control Using Free-Vortex Wake Models*, *Wind Energy Science* (2024). doi: 10.5194/wes-9-721-2024
3. *M. Becker*, B. Ritter, B. Doekemeijer, D. van der Hoek, U. Konigorski, D. Allaerts, and J.W. van Wingerden, *The Revised FLORIDyn Model: Implementation of Heterogeneous Flow and the Gaussian Wake*, *Wind Energy Science* (2022). doi: 10.5194/wes-7-2163-2022.
2. *M. Becker*, D. Allaerts, and J.W. van Wingerden, *Ensemble-Based Flow Field Estimation Using the Dynamic Wind Farm Model FLORIDyn*, *Energies* (2022). doi: 10.3390/en15228589
1. T. Göçmen, F. Campagnolo, T. Duc, I. Eguinoa, S. J. Andersen, V. Petrović, L. Imširović, R. Braunbehrens, J. Liew, M. Baungaard, M. P. Van Der Laan, G. Qian, M. Aparicio-Sanchez, R. González-Lope, V. V. Dighe, *M. Becker*, M. J. van den Broek, J.W. van Wingerden, A. Stock, M. Cole, R. Ruisi, E. Bossanyi, N. Requate, S. Strnad, J. Schmidt, L. Vollmer, I. Sood, and J. Meyers, *FarmConnors Wind Farm Flow Control Benchmark -- Part 1: Blind Test Results*, *Wind Energy Science* (2022). doi: 10.5194/wes-7-1791-2022

CONFERENCE PUBLICATIONS

7. L. Starink, Z. Xie, *M. Becker*, D. van den Berg, and J.W. van Wingerden, *Optimal (Dynamic) Turbine Repositioning Strategies for a Floating Wind Farm Depending on Mooring Line Stiffness*, submitted to the Deep Wind conference proceedings (2025)
6. *M. Becker*, D. Allaerts, and J.W. van Wingerden, *Time-Shifted Cost Function Design for More Efficient Dynamic Wind Farm Flow Control*, 2024 IEEE Conference on Control Technology and Applications (CCTA) (2024). doi: 10.1109/CCTA60707.2024.10666535
5. *M. Becker*, D. Allaerts, and J.W. van Wingerden, *Wind Pattern Clustering of High Frequent Field Measurements for Dynamic Wind Farm Flow Control*, *Journal of Physics: Conference Series* (2024). doi: 0.1088/1742-6596/2767/3/032028
4. E. Taschner, *M. Becker*, R. Verzijlbergh, and J.W. van Wingerden, *Comparison of Helix and Wake Steering Control for Varying Turbine Spacing and Wind Direction*, *Journal of Physics: Conference Series* (2024). doi: 10.1088/1742-6596/2767/3/032023
3. J. Gutknecht, *M. Becker*, E. Taschner, S. Stipa, D. Allaerts, A. Viré, and J.W. van Wingerden, *Active Cluster Wake Mixing*, *Journal of Physics: Conference Series* (2024). doi: 0.1088/1742-6596/2767/9/092052

2. J. Gutknecht, *M. Becker*, C. Muscari, T. Lutz, and J.W. van Wingerden, *Scaling DMD Modes for Modeling Dynamic Induction Control Wakes in Various Wind Speeds*, 2023 IEEE Conference on Control Technology and Applications (CCTA) (2023). doi: 10.1109/CCTA54093.2023.10252400
1. *M. Becker*, D. Allaerts, and J. W. van Wingerden, *FLORIDyn - A Dynamic and Flexible Framework for Real-Time Wind Farm Control*, Journal of Physics: Conference Series (2022). doi: 10.1088/1742-6596/2265/3/032103

PATENTS

1. *M. Becker*, J. Gutknecht, E. Taschner, and J.W. van Wingerden, *Wind turbine control for cluster wake mixing* (2025), WO 2025/051959 A1

