

POLITECNICO DI TORINO

Master of Science in Aerospace Engineering



**Politecnico
di Torino**



**von KARMAN INSTITUTE
FOR FLUID DYNAMICS**

Master's Thesis

Large-eddy simulation of wind farms and the impact of floating solar PV on wind power extraction

Supervisors

Prof. Michele IOVIENO

Prof. Jeroen VAN BEECK

Prof. Wim MUNTERS

Dr. Simone GREMMO

Dr. Baris KALE

PhD Lorenzo SCHENA

PhD Alexandros PALATOS-PLEXIDAS

Candidate

Benito DELLO IACONO

October 2024

Abstract

To address the limitations of expanding offshore wind farms in Belgium, the strategy of integrating floating solar photovoltaic (PV) systems and wind farms has emerged as a viable solution to increase power generation while using the same sea surface. In this study, large-eddy simulations (LES) of wind farms are performed through the ExaWind open-source codes, AMR-Wind and OpenFAST. First, the numerical framework is tested against experimental data and other LES solvers within the SWiFT Neutral Benchmark. Strong agreement is noted in horizontal wind speed and turbulence intensity, indicating that the atmospheric boundary layer is accurately captured. However, generator power and torque exhibit greater variability, with a tendency to overestimate compared to experimental measurements. Nevertheless, all simulated quantities remain within one standard deviation of the experimental data. The wake analysis further shows that the wake strength and shape are well reproduced, though the wake recovery is underestimated. These trends have been previously highlighted with similar solvers for the same benchmark, making the results consistent with the literature. Finally, the validated methodology is applied to assess the impact of floating solar PV panels, represented via an immersed boundary method, on the power extraction of a finite, aligned reference wind farm. A sensitivity analysis on the panels' layout reveals that their presence leads to an overall reduction in wind farm power output, primarily due to a slower wake recovery, which results from a weaker kinetic energy entrainment, and a higher velocity deficit.

Contents

List of Tables	III
List of Figures	IV
1 Introduction	1
2 Numerical framework	6
2.1 LES modeling of ABL flows in AMR-Wind	7
2.2 Turbine implementation in AMR-Wind/OpenFAST	10
2.3 AMR-Wind scaling performance on VKI Ealin43	12
2.3.1 Strong scaling	12
2.3.2 Weak scaling	15
3 A test case: the SWiFT Neutral Benchmark	18
3.1 SWiFT measurements	19
3.2 Simulation methodology	20
3.2.1 Precursor simulation	22
3.2.2 Turbine simulations	23
3.3 Results	26
3.3.1 Atmospheric inflow	26
3.3.2 Wind turbine response	31
3.3.3 Wind turbine wake	36
4 Impact of floating solar PV on wind farm flow and power output	42
4.1 Benchmark wind farm simulation	43
4.1.1 Simulation setup	44
4.1.2 Precursor simulation results	47
4.1.3 Wind farm simulation results	50
4.2 Sensitivity analysis on floating solar PV layouts	56
4.2.1 Power budget analysis	57
4.2.2 Flow analysis	60

5 Conclusions	66
Bibliography	69

List of Tables

2.1	AMR-Wind strong-scaling performance on VKI Ealin43 CPUs. . . .	13
2.2	AMR-Wind weak-scaling test sets.	15
2.3	AMR-Wind weak-scaling performance on VKI Ealin43 CPUs. . . .	16
3.1	Summary of the six data bins collected during the SWiFT experimental campaign and used for validating numerical simulations. The final row of the table presents the ensemble average across all bins.	20
3.2	Description of turbine simulations. All cases are run with AMR-Wind/OpenFAST, as detailed in Section 2.2.	24
3.3	Cost estimates for turbine simulations. The cost is calculated by multiplying the computation hours by the number of cores allocated, according to [50].	25
3.4	Mean \pm standard deviation for time series of WTG _{a1} response quantities of interest.	35
4.1	Overview of the cases performed within the sensitivity study. Here, h , l and w are the height, length and width of the FSPV platform, respectively. The distance between the platform center and the wind turbine is denoted by x_c , and N_{PV} indicates the number of platforms. P/P_0 is the ratio of the wind farm power output with FSPV panels (P) to that of the wind farm without FSPV panels (P_0).	57

List of Figures

1.1	EU’s electricity mix to 2050. Figure adapted from Ref. [5].	2
1.2	Offshore wind energy in the Belgian North Sea. Figure reprinted from Ref. [12].	3
1.3	SeaVolt vision. Figure reprinted from Ref. [16].	4
2.1	Example AMR-Wind grid with two levels of patch-based refinement. Figure reprinted from Ref. [34].	8
2.2	OpenFAST schematic. Figure reprinted from Ref. [40].	11
2.3	Average Walltime per step vs. number of cores for strong scaling. The dashed black line represents the ideal scaling.	14
2.4	Speedup and efficiency vs. number of cores for strong scaling. The dashed black line represents the ideal speedup, while the dotted line shows the best curve that fits the data following Amdahl’s law. . . .	14
2.5	Average Walltime per step vs. number of cores for weak scaling. The dashed black line represents the ideal scaling.	16
2.6	Speedup and efficiency vs. number of cores for weak scaling. The dashed black line represents the ideal speedup, while the dotted line shows the best curve that fits the data following Gustafson’s law. We exclude data for $N > 16$ from the fitting, as the speedup departs importantly from a linear pattern.	17
3.1	SWiFT facility. Figure reprinted from Ref. [46].	19
3.2	Schematic of the SWiFT benchmark setup. Figure reprinted from Ref. [45].	21
3.3	Precursor simulation domain. The red arrow points to the North, while the black and blue one indicates the streamwise direction. . .	22
3.4	Refinement levels for <i>ADM coarse</i> turbine simulation.	24
3.5	Refinement levels for <i>ADM fine</i> and <i>ALM fine</i> turbine simulations.	25
3.6	Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the precursor simulation.	27

3.7	Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.	27
3.8	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.	28
3.9	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.	28
3.10	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [16,000; 20,000]$ s. The orange line is reconstructed from the study by Hsieh et al. (2021) [47]. The black squares are the measurements collected by METa1 meteorological tower [53]. The dashed black line indicates the rotor hub-height.	29
3.11	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, averaged over the time window $t = [16,000; 20,000]$ s. The orange line is reconstructed from the study by Hsieh et al. (2021) [47]. The dashed black line indicates the rotor hub-height.	30
3.12	Time series of hub-height wind speed, U_{HH} . The blue line represents the output from AMR-Wind, while the gray line is an overlay of the six individual experimental time series recorded by METa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.	30
3.13	Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with <i>ADM coarse</i> approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.	32
3.14	Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with <i>ADM fine</i> approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.	33

3.15	Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with <i>ALM fine</i> approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.	34
3.16	Mean (colored bar) and standard deviation (horizontal black line) for time series of wind turbine response quantities.	35
3.17	Time-averaged normalized velocity deficit in YZ-planes located at $x/D = 2, 3, 4$ and extending 1D from the hub center in both positive and negative lateral and vertical direction. Results are shown for <i>ADM coarse</i> , <i>ADM coarse</i> and <i>ADM coarse</i> cases. The black circle outlines rotor area.	36
3.18	Resolved turbulence intensity in YZ-planes located at $x/D = 2, 3, 4$ and extending 1D from the hub center in both positive and negative lateral and vertical direction. Results are shown for <i>ADM coarse</i> , <i>ADM coarse</i> and <i>ADM coarse</i> cases. The black circle outlines rotor area.	37
3.19	Horizontal profiles of normalized velocity deficit (a) and resolved turbulence intensity (b) averaged in time and sampled at hub height for different downstream distances. The black dashed lines indicate the rotor edges.	38
3.20	Vertical profiles of normalized velocity deficit (a) and resolved turbulence intensity (b) averaged in time and sampled at hub height for different downstream distances. The black dashed lines indicate the rotor edges.	39
3.21	Maximum of time-averaged velocity deficit fields as a function of the downstream position. PALM, Nalu-Wind and SpinnerLidar data are taken from [51].	40
4.1	Benchmark wind farm layout.	44
4.2	IEA Wind 15-MW offshore reference wind turbine. Figure reprinted from Ref. [66].	45
4.3	Refinement level for wind farm simulation.	46
4.4	Turbine controller regulation trajectory. Figure adapted from Ref. [66].	46
4.5	Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the precursor simulation.	47
4.6	Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.	47

4.7	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.	48
4.8	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.	48
4.9	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [19,200; 21,600]$ s. The dashed black line indicates the rotor hub-height.	49
4.10	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, averaged over the time window $t = [19,200; 21,600]$ s. The dashed black line indicates the rotor hub-height.	49
4.11	Streamwise velocity field in the turbine hub-height plane at $t = 21,600$ s. The dashed black lines divide the plane into four equal bands across which the spatial average is computed.	50
4.12	Streamwise velocity field in the symmetry XZ-plane of the third column turbines at $t = 21,600$ s.	50
4.13	Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the wind farm simulation.	51
4.14	Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.	51
4.15	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the wind farm simulation. The dashed black line indicates the rotor hub-height.	52
4.16	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the wind farm simulation. The dashed black line indicates the rotor hub-height.	52
4.17	Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [20,400; 21,600]$ s. The dashed black line indicates the rotor hub-height.	53
4.18	Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, averaged over the time window $t = [20,400; 21,600]$ s. The dashed black line indicates the rotor hub-height.	53
4.19	Normalized streamwise velocity deficit averaged over time and across the $7D$ -wide XY-strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.	54
4.20	Normalized streamwise velocity deficit averaged over time and across the symmetry XZ-planes of the turbine columns (indicated by solid black lines).	54

4.21	Normalized turbulent kinetic energy averaged across the $7D$ -wide XY-strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.	55
4.22	Normalized turbulent kinetic energy averaged across the symmetry XZ-planes of the turbine columns (indicated by solid black lines).	55
4.23	Normalized streamwise turbulence intensity averaged across the $7D$ -wide XY-strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.	55
4.24	Normalized streamwise turbulence intensity averaged across the symmetry XZ-planes of the turbine columns (indicated by solid black lines).	55
4.25	Normalized kinematic lateral turbulent momentum flux averaged across the $7D$ -wide XY-strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.	56
4.26	Normalized kinematic vertical turbulent momentum flux averaged across the symmetry XZ-planes of the turbine columns (indicated by solid black lines).	56
4.27	Normalized power output P/P_0 as a function of the wind farm rows for all cases analyzed within the sensitivity study.	58
4.28	Normalized power output P/P_0 as a function of the wind farm rows for different values of the h/z_h parameter.	58
4.29	Normalized power output P/P_0 as a function of the wind farm rows for different values of the l/D parameter.	59
4.30	Normalized power output P/P_0 as a function of the wind farm rows for different values of the x_c/D parameter.	59
4.31	Time-averaged vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, for all cases. The dashed black line indicates the rotor hub-height.	60
4.32	Time-averaged vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, for all cases. The dashed black line indicates the rotor hub-height.	61
4.33	Side view (top) and top view (bottom) of FSPV layout in Case 2.	61
4.34	Side view (top) and top view (bottom) of FSPV layout in Case 7.	62
4.35	Delta of normalized streamwise velocity deficit in the hub-height XY-plane, obtained by subtracting the Case 0 field from the Case 2 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case.	62

4.36	Delta of normalized streamwise velocity deficit in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case. Light brown areas indicate out-of-scale values.	63
4.37	Delta of normalized turbulent kinetic energy in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Light brown areas indicate out-of-scale values.	63
4.38	Delta of normalized kinematic vertical turbulent momentum flux in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Light brown areas indicate out-of-scale values.	63
4.39	Delta of normalized streamwise velocity deficit in the hub-height XY-plane, obtained by subtracting the Case 0 field from the Case 7 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case.	64
4.40	Delta of normalized streamwise velocity deficit in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case. Light brown areas indicate out-of-scale values.	64
4.41	Delta of normalized turbulent kinetic energy in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field.	64
4.42	Delta of normalized kinematic vertical turbulent momentum flux in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field.	64

Chapter 1

Introduction

Climate change represents one of the most urgent and complex challenges of our time. The effects of global warming, already evident in rising global temperatures, increasing sea levels, and more frequent extreme weather events, demand decisive and coordinated action worldwide. The Paris Agreement, adopted in 2015 by 196 countries, is a cornerstone in the global fight against climate change, setting the goal of limiting global temperature rise to well below 2°C above pre-industrial levels, with the ambition to restrict it to 1.5°C [1]. To achieve this goal, a rapid transition towards low-carbon energy sources is imperative. The global energy mix must evolve quickly to reduce reliance on fossil fuels and increase the share of energy generated from renewable sources. According to the International Renewable Energy Agency (IRENA), in order to keep global warming within the 1.5°C limit, the share of renewable energy in the global energy mix should rise from 16% in 2020 to approximately 77% by 2050 [2].

In this context, wind energy is not only one of the lowest-cost electricity sources per unit of energy produced, but it also shares the lowest mean life-cycle CO₂-equivalent emissions together with nuclear energy, placing it among the cleanest energy sources [3]. These factors have contributed to the exponential growth of wind energy in recent years. Globally, installed wind capacity passed the symbolic milestone of 1 TW by the end of 2023 and, at the present pace, it is expected to reach 2 TW before 2030 [4]. In the EU-27, wind energy will represent about 50% of the electricity mix as projected by WindEurope [5] (Figure 1.1).

Global offshore wind capacity currently account for only 9% of the total wind power installation [4]. However, stronger and more consistent winds, along with the ability to utilize vast marine areas without the visual and noise impacts associated with land-based installations, give offshore wind energy immense potential. In addition, the advent of floating wind systems allows access to sea areas where water depth exceeds 60 meters, which represent about 80% of the global offshore wind resource potential [6]. As a result, the number of offshore installations is expected

to roughly double by 2030 [7]. In Europe, only 13% of total wind power capacity is currently installed offshore, although offshore installations are projected to reach one-third of the new wind installations by 2030 [8].




Figure 1.1: EU’s electricity mix to 2050. Figure adapted from Ref. [5].

Belgium had the sixth-highest offshore wind capacity globally in 2021 according to International Energy Agency (IEA). In 2004, a 156-square kilometer area of the North Sea within Belgium’s Exclusive Economic Zone (EEZ) was designated for wind farm development. Eight wind farms have been allocated in this region, the first of which became operational in 2009. In December 2020, the installed offshore wind capacity amounts to 2,262 MW, which corresponds to the annual electricity needs of nearly 2 million households, or about 10% of the country’s total electricity demand [9]. Moreover, a second more western area, the Princess Elisabeth Zone, has already been assigned with the goal of increasing this capacity to 5.8 GW [10] (Figure 1.2). Nevertheless, further development beyond this second zone is expected to be complex due to expansion constraints and wake effects. Consequently, Belgium is exploring various options to optimize offshore capacity within its EEZ, including the investment in floating solar photovoltaic (FSPV) panels, which could supply an additional 1 GW [11].

Offshore solar and wind energy are highly complementary. Wind power generation is typically higher at night and during winter when sunlight is limited or absent. In contrast, solar power output peaks during the day and in summer. Additionally, by adding offshore solar in between the sea surface covered by wind




Figure 1.2: Offshore wind energy in the Belgian North Sea. Figure reprinted from Ref. [12].

turbines, up to five times more renewable power can be extracted in the same area [13], helping to address space limitations.

In September 2023, the SeaVolt Consortium [14] successfully installed its first floating test platform, consisting of eight FSPV panels, in the Blue Accelerator test zone near the port of Ostend, Belgium. Thanks to this installation, the consortium is gathering crucial data to advance the technology and aims to deploy the first full-scale platforms between 2025 and 2026 (Figure 1.3). Similarly, in May 2024, the Dutch offshore solar company Oceans of Energy [15] completed the North Sea Three (NS3) floating solar farm in the same port.

Many FSPV panels could significantly change the wind flow aloft that hits the wind turbines. First, they can be seen as a step change in surface roughness, creating an internal boundary layer (IBL) that reduces wind speed and increases turbulence. Furthermore, the surface heating of the panels adds further turbulence to the wind flow. These effects could potentially alter the performance of a wind farm, particularly in terms of power extraction and turbine loading. The scope of this work is to study the impact of solar PV on wind farm flow and power extraction using large eddy simulations (LES).

In industrial applications where statistics on mean flow and power are needed, engineering models using Reynolds-averaged Navier-Stokes (RANS) equations are well suited. However, to fully understand wind farm aerodynamics within the




Figure 1.3: SeaVolt vision. Figure reprinted from Ref. [16].

atmospheric boundary layer (ABL), a more advanced approach is necessary to resolve, rather than model, at least the significant eddies that influence wake evolution and power generation. Direct numerical simulations (DNS), which resolve all flow eddies, offer the most comprehensive solution but are currently impractical due to limitations in computing resources. By resolving the largest eddies and modeling only the smallest scales, LES offer a more viable solution for high-fidelity modeling of wind farms. An extensive review of LES for wind farm aerodynamics can be found in Mehta et al. (2014) [17] and Breton et al. (2017) [18].

This manuscript is part of the work carried out during a five-month internship at the von Karman Institute for Fluid Dynamics (VKI) and is structured as follows.

Chapter 2 provides a detailed description of the numerical framework employed in this work, which includes the LES solver AMR-Wind and the whole-turbine simulation tool OpenFAST, both part of the ExaWind software suite. Additionally, AMR-Wind strong and weak scaling tests are performed for an ABL simulation on the VKI workstation Ealin43.

Chapter 3 focuses on testing AMR-Wind, also coupled with OpenFAST, within the SWiFT benchmark with neutral atmospheric stratification, which involves a small-scale test case with two wind turbines. This benchmark offers extensive experimental and numerical data and has been widely used in recent years to validate similar LES solvers. We examine atmospheric inflow, turbine response and wake using both the actuator disk model (ADM) and the actuator line model (ALM) with varying grid resolutions. The results are compared with in situ measurements and reference literature.

In Chapter 4, the methodology developed and tested in Chapter 3 is applied to assess the impact of FSPV panels on the power extraction of a finite benchmark wind farm. The panels are implemented using the immersed boundary method

(IBM). A sensitivity analysis on panels' layout is conducted, analyzing the variations in the wind farm's power output and flow statistics. Furthermore, we evaluate the potential for the panels to enhance the power generation of the wind farm, as shown in previous studies on windbreaks, buildings, or hills.

Finally, Chapter 5 presents a comprehensive summary of the main findings of this thesis.

Chapter 2

Numerical framework

In the present work, we rely on the open-source suite of codes ExaWind [19], developed as part of the Exascale Computing Project funded by the U.S. Department of Energy. This project aims to achieve the highest-fidelity simulations of wind turbines and wind farms to date, enhancing the understanding of crucial phenomena that affect the wind plant performance, including wake formation, complex-terrain impacts, wake-atmosphere interaction, turbine-turbine interaction, and blade-boundary-layer dynamics. By accurately capturing these complex physical phenomena, high-fidelity models play a crucial role in advancing the design and operation of wind energy systems. Moreover, these models act as benchmarks for developing next-generation low-fidelity models and the extensive data they generate represent a valuable resource for informing upcoming machine-learning-based engineering models. Additionally, with ongoing advancements in high-performance computing (HPC), a key objective of the ExaWind codes is the ability to scale efficiently on modern supercomputers, including those accelerated with graphical-processing units (GPUs).

ExaWind consists of three physics-based codes: AMR-Wind [20] and Nalu-Wind [21], which are two incompressible fluid dynamics solvers for wind turbine and wind farm simulations; and OpenFAST [22], a whole-turbine simulation tool. Together, these codes have the capability to perform simulations across a wide range of fidelity. Specifically, AMR-Wind can perform LES of wind farms in ABL flows using ADM [23] or ALM [24] for the turbines. Alternatively, it can serve as a background solver when coupled with Nalu-Wind for higher-fidelity blade-resolved simulations.

AMR-Wind has been demonstrated to be five times faster than Nalu-Wind for ABL simulations [25]. In addition, blade-resolved simulations of wind turbines and wind farms would be impracticable for our computational resources due to the spatial resolution required to resolve the boundary layer over the blades. Therefore, in this study, we focus exclusively on running cases with actuator models, ADM or ALM, coupling AMR-Wind with OpenFAST. A complete description of ADM and

ALM is available in Tossas and Leonardi (2013) [26].

2.1 LES modeling of ABL flows in AMR-Wind

AMR-Wind is a relatively newborn open-source CFD code written in C++ and developed by a team of experts from Lawrence Berkeley National Laboratory (LBNL), National Renewable Energy Laboratory (NREL), and Sandia National Laboratories (SNL). In 2022, Min et al. (2022) [27] investigated AMR-Wind performance for the LES simulation of ABL flows on Oak Ridge Leadership Facility supercomputers Summit [28] and Crusher, the testbed for the Frontier exascale system [29], performing strong and weak scaling tests. Chatterjee et al. (2022) [30] used AMR-Wind to study the interaction between atmospheric turbulence and wind farms operating in coastal low-level jet events with mesoscale-driven LES. In 2023, Cheung et al. (2023) [31] employed AMR-Wind with OpenFAST coupled ADM to conduct LES of five wind farms in unstable atmosphere, matching the numerical results with measurements gathered in the American WAKE experimeNt (AWAKEN) [32], a multi-institutional field campaign designed to address ABL-wind-farm interactions.

The code functions as a block-structured incompressible flow solver with adaptive mesh refinement (AMR) capabilities, built on top of the AMReX library [33]. This library provides the necessary grid data structures, adaptivity features, and linear solvers for solving flow governing equations. Figure 2.1 shows an example of AMR-Wind grid created using AMReX.

AMR-Wind solves the incompressible Navier-Stokes equations with variable density and viscosity, along with the scalar transport equation for potential temperature. Here, the equations for continuity, momentum and potential temperature are expressed in the LES formulation with the Boussinesq approximation [35], the Einstein notation and in non-dimensional form:

$$\frac{\partial \tilde{u}_j}{\partial x_j} = 0, \quad (2.1)$$

$$\frac{\partial \tilde{u}_i}{\partial t} + \tilde{u}_j \frac{\partial \tilde{u}_i}{\partial x_j} = -\frac{1}{\tilde{\rho}} \frac{\partial \tilde{p}}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j} - \frac{\tilde{\theta} - \theta_0}{\theta_0} g_i - f \varepsilon_{ij3} \tilde{u}_j + f_i, \quad (2.2)$$

$$\frac{\partial \tilde{\theta}}{\partial t} + \tilde{u}_j \frac{\partial \tilde{\theta}}{\partial x_j} = -\frac{\partial \tau_{\theta j}}{\partial x_j}, \quad (2.3)$$

where the tilde denotes the resolved-scale component of the quantities obtained after the LES filtering operation, while we indicate the subgrid-scale (SGS) part with double primes. \tilde{u}_i represents the i th component of the velocity vector, $\tilde{\rho}$ the density, \tilde{p} the pressure, g_i the gravitational acceleration vector, which is nonzero




Figure 2.1: Example AMR-Wind grid with two levels of patch-based refinement. Figure reprinted from Ref. [34].

only in the vertical direction (z), f the Coriolis parameter and f_i corresponds to any additional source terms (e.g. body force). In the term $(\theta - \theta_0)/\theta_0$, which governs the buoyancy force, θ is the potential temperature and θ_0 is the reference potential temperature. τ_{ij} and $\tau_{\theta j}$ in the momentum and energy equations are the shear stress tensor and the heat flux vector respectively, and are defined as

$$\tau_{ij} = -\frac{2}{Re} S_{ij} + \tau_{ij}'' = -\frac{1}{Re} \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right) + \tau_{ij}'', \quad (2.4)$$

$$\tau_{\theta j} = -\frac{1}{Pe} \frac{\partial \tilde{\theta}}{\partial x_j} + \tau_{\theta j}'', \quad (2.5)$$

where Re is the Reynolds number, Pe is the Peclet number, S_{ij} is the resolved-scale strain-rate tensor, and τ_{ij}'' and $\tau_{\theta j}''$ are the SGS shear stress tensor and heat flux vector, respectively. Since the typical grid spacing in LES of ABL (\sim m) is several orders of magnitude larger than the size of smallest turbulent scales (\sim mm), the resolved part of τ_{ij} and $\tau_{\theta j}$ is practically negligible and they are dominated by the SGS terms.

The system of filtered equations is closed only once the SGS fluxes, τ_{ij}'' and $\tau_{\theta j}''$, are determined as a function of the resolved-scale field. For the SGS modeling, AMR-Wind employs the constant coefficients model from Moeng (1984) [36], which

solves one prognostic equation for the subgrid-scale turbulent kinetic energy (SGS-TKE), k'' :

$$\frac{\partial k''}{\partial t} = -\tilde{u}_j \frac{\partial k''}{\partial x_j} - \widetilde{u_i'' u_j''} \frac{\partial \tilde{u}_i}{\partial x_j} + \frac{g}{\theta_0} \widetilde{w'' \theta''} - \frac{\partial [\widetilde{u_j'' k''} + \widetilde{u_j'' p''} / \rho_0]}{\partial x_j} - \epsilon, \quad (2.6)$$

where w is the wind speed component in the vertical direction and ρ_0 is the reference density. The terms on the right-hand side represent, in order, advection by the resolved-scale motion, local shear production, local buoyancy production or consumption, transport by SGS motion and pressure, and the dissipation rate of SGS-TKE. The SGS fluxes are related to the resolved field through the Boussinesq eddy viscosity hypothesis [37] as

$$\tau_{ij}'' = -2K_M S_{ij} = -K_M \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right), \quad (2.7)$$

$$\tau_{\theta j}'' = -K_H \frac{\partial \tilde{\theta}}{\partial x_j}, \quad (2.8)$$

where K_M and K_H are the SGS eddy coefficients for momentum and heat and are computed from k'' . Furthermore, we remark that the LES filter is not explicitly defined, but instead implicitly determined by the grid spacing, as with almost all LES models for ABL flows.

The vertical SGS fluxes are also modeled at the terrain boundary following Moeng (1984) [36]:

$$(\tau_{xz}'')_0 = \langle \tau_{xz}'' \rangle_0 \frac{S_1 \langle \tilde{u}_1 \rangle + \langle S_1 \rangle (\tilde{u}_1 - \langle \tilde{u}_1 \rangle)}{\langle S_1 \rangle \langle \tilde{u}_1 \rangle}, \quad (2.9)$$

$$(\tau_{yz}'')_0 = \langle \tau_{yz}'' \rangle_0 \frac{S_1 \langle \tilde{v}_1 \rangle + \langle S_1 \rangle (\tilde{v}_1 - \langle \tilde{v}_1 \rangle)}{\langle S_1 \rangle \langle \tilde{v}_1 \rangle}, \quad (2.10)$$

$$(\tau_{\theta z}'')_0 = \langle \tau_{\theta z}'' \rangle_0 \frac{S_1 (\langle \tilde{\theta}_1 \rangle - \theta_0) + \langle S_1 \rangle (\tilde{\theta}_1 - \langle \tilde{\theta}_1 \rangle)}{\langle S_1 \rangle (\langle \tilde{\theta}_1 \rangle - \theta_0)} \quad (2.11)$$

where the subscript 0 refers to the values at the surface, subscript 1 the values at the first cell height in the vertical direction and $\langle \cdot \rangle$ represents the horizontal plane average. $S = \sqrt{\tilde{u}^2 + \tilde{v}^2}$ is the magnitude of horizontal wind speed, where \tilde{u} and \tilde{v} are the resolved horizontal wind speed components in longitudinal and lateral directions, x and y , respectively. The mean SGS stresses $\langle \tau_{xz}'' \rangle_0$, $\langle \tau_{yz}'' \rangle_0$ and $\langle \tau_{\theta z}'' \rangle_0$ are computed using the Monin-Obukhov similarity law [38].

For simplicity, we omit the tilde notation for the rest of the work. We also always use $\bar{\cdot}$ to represent time-averages and $\langle \cdot \rangle$ to denote spatial averaging in the horizontal directions.

Discretization in AMR-Wind is performed with a spatially and temporally second-order accurate finite-volume scheme. It adopts a partial staggering for spatial discretization, computing velocity, scalar quantities, and pressure gradients at cell centers, while pressure is determined at nodes. Staggering is also used for time discretization, similar to a Crank–Nicholson formulation:

$$\frac{u_i^* - u_i^n}{\Delta t} + \left[\frac{\partial u_i u_j}{\partial x_j} \right]^{n+1/2} = \frac{1}{\rho^{n+1/2}} \left(\frac{\partial \tau_{ij}^{n+1}}{\partial x_j} - \frac{\partial p^{n-1/2}}{\partial x_i} \right) + F_i^{n+1/2}, \quad (2.12)$$

$$\frac{c_k^{n+1} - c_k^n}{\Delta t} + \left[\frac{\partial c_k u_j}{\partial x_j} \right]^{n+1/2} = \frac{1}{\rho^{n+1/2}} \frac{\partial q_{kj}^{n+1}}{\partial x_j} + G_k^{n+1/2}, \quad (2.13)$$

where u is velocity, c indicates a scalar quantity and q its diffusive transport flux. The index n refers to the time step, i runs over the momentum equations, k runs over the scalar equations, and j denotes summation. $F_i^{n+1/2}$ and $G_k^{n+1/2}$ are source terms, evaluated at the time step $n + 1/2$. In variable-density simulations, the density ρ is calculated at timestep $n + 1/2$ by a simple average:

$$\rho^{n+1/2} = \frac{1}{2}(\rho^n + \rho^{n+1}). \quad (2.14)$$

The time stepping scheme employs a predictor–corrector approach [39]. In the predictor step, a lagged pressure is used to predict an intermediate velocity, u_i^* , which is not guaranteed to be divergence free. The corrector step applies a projection to enforce the divergence-free constraint on the velocity field and updates the pressure.

By using partial staggering and an approximate projection method, AMR-Wind constructs linear systems that can be solved using standard techniques like multilevel multigrid (MLMG), achieving a good measure of both efficiency and accuracy. For more details on discretization in AMR-Wind the reader is referred to work by Almgren et al. (1998) [39].

2.2 Turbine implementation in AMR-Wind/ OpenFAST

The wind turbines in AMR-Wind are represented using actuator disk and actuator line models, which can be fully integrated into the code (e.g., Joukowski actuator disk) or coupled with OpenFAST. Typically, the former are less accurate but more computationally efficient.

OpenFAST is an open-source multi-physics, multi-fidelity wind turbine simulation tool mostly written in Fortran 2003 and maintained by a dedicated team at the NREL. OpenFAST “couples computational modules for aerodynamics, hydrodynamics for offshore structures, control and electrical system (servo) dynamics, and

structural dynamics to enable coupled nonlinear aero-hydro-servo-elastic simulation in the time domain. OpenFAST enables the analysis of a range of wind turbine configurations, including two- or three-blade horizontal-axis rotor, pitch or stall regulation, rigid or teetering hub, upwind or downwind rotor, and lattice or tubular tower. The wind turbine can be modeled on land or offshore on fixed-bottom or floating substructures” [40]. The schematic of OpenFAST is presented in Figure 2.2.




Figure 2.2: OpenFAST schematic. Figure reprinted from Ref. [40].

When the actuator methods in AMR-Wind are coupled with OpenFAST, the two codes operate as follows: first, AMR-Wind completes a fluid solve, performs a local search to identify the ADM/ALM point locations at that time step, and then interpolates and passes the fluid information at the actuator points to OpenFAST. Next, OpenFAST iterates in time while keeping the fluid domain properties fixed, solving for the new forces and locations of the actuator points at the next fluid timestep to pass to AMR-Wind. OpenFAST also computes and outputs other turbine-specific quantities of interest, such as rotor speed, generator torque, and power. AMR-Wind then updates the actuator point locations and applies the actuator forces in its next momentum solve as a body force, f_i :

$$f_i = \int_0^L F_i(l)g(\mathbf{r}) dl, \quad (2.15)$$

where l is the distance along an actuator element, F_i denotes the actuator element forces computed from the incoming velocity flow and the airfoil characteristics specified in OpenFAST, and $g(\mathbf{r})$ is an isotropic Gaussian kernel that spreads the actuator forcing to the fluid domain:

$$g(\mathbf{r}) = \frac{1}{\varepsilon^3 \pi^{3/2}} e^{-\left(\frac{|\mathbf{r}|}{\varepsilon}\right)^2}, \quad (2.16)$$

where \mathbf{r} represents the distance from the center of the actuator element to the point where the body force is imposed and ε is a characteristic length scale that determines the volume over which the body forces are spread. Coupling is done with fixed timesteps in both codes to allow for precise subsampling by OpenFAST.

2.3 AMR-Wind scaling performance on VKI Ealin43

Here, we discuss the strong and weak scaling performance of AMR-Wind on the VKI workstation Ealin43, which was used to conduct all the simulations documented in this work. The machine is equipped with 2 AMD 9654 Genoa 96-core CPUs, 2 NVIDIA RTX A5000 and 1 NVIDIA A100 GPUs. However, we consider CPU-only performance, as GPU acceleration was neither explored nor used in this study.

2.3.1 Strong scaling

The strong scaling performance is evaluated for a reference ABL simulation on a fixed $4,000 \text{ m} \times 4,000 \text{ m} \times 960 \text{ m}$ domain with a uniform mesh resolution $\Delta x = 10 \text{ m}$. The timestep size is kept constant at $\Delta t = 0.5 \text{ s}$ as well. In case of strong scaling, we increase the number of cores allocated while the problem size (domain size, spatial and temporal spacings) remains constant, resulting in a reduction of the workload per processor (cells/core).

Table 2.1 summarizes the strong-scaling performance, in terms of average wall-time per step, speedup and efficiency. The speedup is usually defined as

$$\text{Speedup} = t(1)/t(N), \quad (2.17)$$

where $t(1)$ is the walltime required to complete a run on a single processor and $t(N)$ the time to complete the same run on N processors. In the following, the terms *processor* and *core* are used interchangeably. In our case, a *run* refers to the duration of an individual timestep.

The efficiency is defined as the actual speedup divided by the ideal speedup, which is equal to the number of cores (N):

$$\text{Efficiency (\%)} = \frac{\text{Speedup}}{N} \times 100. \quad (2.18)$$

Figure 2.3 illustrates the typical time vs. cores plot. The walltime per each step scales almost ideally up to 16 cores and then it flattens out as the number of processors increases. The same trend can be observed looking at the speedup

Table 2.1: AMR-Wind strong-scaling performance on VKI Ealin43 CPUs.

Cores	Cells/core	Walltime per Δt [s]	Speedup [-]	Efficiency (%)
1	1.54e7	77.61	1.00	100.00
2	7.68e6	38.59	2.01	100.54
4	3.84e6	19.77	3.92	98.12
8	1.92e6	9.94	7.81	97.62
16	9.60e5	5.12	15.16	94.73
32	4.80e5	2.89	26.87	83.98
64	2.40e5	1.64	47.44	74.12
96	1.60e5	1.37	56.80	59.17
128	1.20e5	1.23	62.99	49.21
160	9.60e4	1.12	69.25	43.28
192	8.00e4	1.08	72.08	37.54

profile (Figure 2.4). The speedup is limited by the serial part of the program $s = 1 - p$ that cannot be parallelized, as Amdahl’s law [41] states:

$$\text{Speedup} = 1/(s + p/N). \tag{2.19}$$

Only if $s = 0$ does the program achieve the ideal speedup, which would be linear and equal to the number of processors (Speedup = N). This means that each processor would be contributing 100% of its computational power. However, perfect scalability is very rarely attained also due to communication overhead, especially at larger process counts. In Figure 2.4 we plot the curve that best fits the speedup data following Amdahl’s law. Based on this fit, we estimate that the parallel fraction of the program is approximately $p = 0.9918$, indicating that around 99.18% of the workload is parallelizable.

In conclusion, increasing the number of cores accelerates the execution of the program (the walltime decreases and, consequently, the speedup increases), but at the same time, the efficiency of the parallelization drops significantly as we move further away from ideal scaling (Figure 2.4).




Figure 2.3: Average Walltime per step vs. number of cores for strong scaling. The dashed black line represents the ideal scaling.




Figure 2.4: Speedup and efficiency vs. number of cores for strong scaling. The dashed black line represents the ideal speedup, while the dotted line shows the best curve that fits the data following Amdahl's law.

2.3.2 Weak scaling

In contrast to strong scaling, in case of weak scaling we increase both the problem size and the number of processors, so that the workload per processor remains constant. The domain size is increased in the x and y directions, while domain height is fixed (Table 2.2). The spatial and temporal resolution are kept constant at $\Delta x = 10$ m and $\Delta t = 0.5$ s, respectively.

Table 2.2: AMR-Wind weak-scaling test sets.

Cores	Domain size [m ³]	Grid points
1	2,000 × 2,000 × 960	200 × 200 × 96
4	4,000 × 4,000 × 960	400 × 400 × 96
16	8,000 × 8,000 × 960	800 × 800 × 96
64	16,000 × 16,000 × 960	1,600 × 1,600 × 96
100	20,000 × 20,000 × 960	2,000 × 2,000 × 96
144	24,000 × 24,000 × 960	2,400 × 2,400 × 96

Table 2.3 provides the weak-scaling performance, in terms of average walltime per step, speedup and efficiency. In weak scaling the efficiency is usually defined as

$$\text{Efficiency} = \frac{t(1)}{t(N)}. \quad (2.20)$$

Consequently, the speedup can be computed as

$$\text{Speedup} = \text{Efficiency} \times N. \quad (2.21)$$

Figure 2.5 and Figure 2.6 show the walltime, speedup and efficiency as function of the number of cores. In a perfectly weak-scaling application the walltime would remain constant, meaning the program solves progressively larger problems on a larger machine in the same amount of time it takes to solve smaller problems on a smaller machine. This holds nearly true only up to 16 cores, after which the wall time rises significantly as the number of processors grows, more than doubling for $N = 144$. Indeed, the speedup deviates substantially from the ideal profile for $N > 16$, as the efficiency drops rapidly.

In Figure 2.6 we also plot the curve that best fits the speedup data for $N = 1, 4, 16$ and follows Gustafson’s law [42], which argues that the parallel portion scales linearly with the available resources, while the serial portion remains constant:

$$\text{Speedup} = s + p \times N. \quad (2.22)$$

We exclude data for core counts higher than 16, as the speedup can no longer be fit by a linear relationship. From this fit we estimate that for $N \leq 16$ the parallel and the serial fractions of the program are approximately $p = 0.915$ and $s = 0.085$, respectively.

Table 2.3: AMR-Wind weak-scaling performance on VKI Ealin43 CPUs.

Cores	Walltime per Δt [s]	Speedup [-]	Efficiency (%)
1	18.99	1.00	100.00
4	19.75	3.85	96.14
16	20.66	14.71	91.91
64	24.87	48.86	76.34
100	31.50	60.28	60.28
144	40.84	66.95	46.49




Figure 2.5: Average Walltime per step vs. number of cores for weak scaling. The dashed black line represents the ideal scaling.




Figure 2.6: Speedup and efficiency vs. number of cores for weak scaling. The dashed black line represents the ideal speedup, while the dotted line shows the best curve that fits the data following Gustafson’s law. We exclude data for $N > 16$ from the fitting, as the speedup departs importantly from a linear pattern.

Chapter 3

A test case: the SWiFT Neutral Benchmark

One of the main objectives of this thesis work is to experiment with AMR-Wind and its coupling with OpenFAST, from the installation of the codes to the simulation of large wind farms. In this chapter we seek to validate the numerical framework presented in Chapter 2.

In the summer of 2017, an experimental wake steering campaign was conducted at the Scaled Wind Farm Technology (SWiFT) facility [43], made possible by a collaborative effort between SNL and NREL, with funding from the Department of Energy’s Wind Energy Technologies Office A2e program. Based on this experiment, three benchmarks were defined within the IEA Wind Task 31 (WakeBench) [44], differing in terms of atmospheric stability [45]:

- Wake Evolution in a Nearly Neutral Atmosphere
- Wake Evolution in a Stable Atmosphere
- Wake Dynamics in an Unstable Atmosphere

Completed in 2021, these benchmarks represent a comprehensive study intended to compare several simulation tools in terms of accuracy and computational cost, while also validating them against experimental data. However, to the best of our knowledge, AMR-Wind has not been tested within the SWiFT benchmarks, unlike Nalu-Wind. This gap provides an opportunity to validate AMR-Wind in comparison with other similar solvers.

Although the experimental data from the wake steering campaign embraces a variety of conditions, this study focuses exclusively on experiments and simulations with near-neutral atmospheric stability conditions of inflow (near sunrise and sunset), which represent the simplest environmental scenario, and zero nominal yaw offset.

3.1 SWiFT measurements

The SWiFT facility, located at Texas Tech University’s National Wind Institute Research Center in Lubbock, Texas, is a state-of-the-art technology accelerator designed to support wind-plant and turbine research development. Situated in the Great Plains, this facility offers a high wind resource, an unobstructed landscape, low turbulence and a distinct diurnal cycle, making it ideal for isolating the effects of atmospheric stratification on wake dynamics. The facility is equipped with three research-scale Vestas V27, each with a rotor diameter (hereafter referred to as D) of 27 m and a hub height of 32.1 m. Two of these turbines (WTGa1 and WTGb1) are spaced $3D$ apart, perpendicular to the prevailing wind direction, while the third (WTGa2) is positioned $5D$ downwind of WTGa1. The site also includes two 60-meter tall meteorological (MET) towers instrumented with three-dimensional sonic anemometers at five different heights, providing detailed inflow and atmospheric data.



Figure 3.1: SWiFT facility. Figure reprinted from Ref. [46].

For the WakeBench campaign inflow, turbine and wake data have been acquired. Inflow data is obtained from the upwind meteorological tower, METa1, located $2.5D$ upstream of WTGa1, while turbine data is collected from the WTGa1 itself. The database consists of six distinct 10-minute bins, which are summarized in Table 3.1 along with the ensemble average of all bins. The quantities of interest listed in the table include mean hub-height wind speed (\bar{U}_{HH}), power-law wind shear exponent (α), resolved hub-height turbulence intensity (TI_R), and generator power. The α parameter is derived from measurements at 10 m and 58 m, which correspond to the locations of the lowest and highest sonic anemometers on the

meteorological tower, respectively:

$$\alpha = \frac{\log(\bar{U}_{58m}/\bar{U}_{10m})}{\log(58/10)}. \quad (3.1)$$

For details on the calculation of TI_R from the measured full range turbulence intensity, the reader is referred to the work by Hsieh et al. (2021) [47].

The wake data have been collected through the use of a rear-facing, nacelle-mounted scanning lidar from the Technical University of Denmark (DTU), the DTU SpinnerLidar. This continuous-wave lidar acquires 984 line-of-sight velocity measurements in a rosette pattern on a spherical surface during a single scan and at a specific downstream distance x . Each scan lasts approximately 2 s. For the neutral benchmark, wake measurements have been performed between $x/D = 2$ and $x/D = 5$, with a spacing of 1D. Further information on the SpinnerLidar measurements at the SWiFT facility is available in earlier publications [48, 49]. A schematic of the instrumentation used in this benchmark is shown in Figure 3.2.

Table 3.1: Summary of the six data bins collected during the SWiFT experimental campaign and used for validating numerical simulations. The final row of the table presents the ensemble average across all bins.

	\bar{U}_{HH} [m/s]	α [-]	TI_R [-]	Power [kW]
Bin 1	8.19	0.13	0.079	68.5
Bin 2	8.97	0.11	0.065	83.7
Bin 3	8.63	0.13	0.057	78.8
Bin 4	8.55	0.17	0.055	79.3
Bin 5	8.76	0.17	0.055	74.5
Bin 6	9.06	0.13	0.054	89.7
En. Av.	8.69	0.14	0.061	79.1

3.2 Simulation methodology

The simulation methodology adopted for this validation study is similar to that proposed by Hsieh et al. (2021) [47]. The approach consists of two main stages. First, a precursor simulation is performed using a relatively coarse and uniform grid along with a larger timestep to achieve converged ABL statistics. This initial phase typically requires a significant simulation time (5-6 hours) to establish a




Figure 3.2: Schematic of the SWiFT benchmark setup. Figure reprinted from Ref. [45].

fully developed turbulent flow field. Following this, AMR-Wind is coupled with OpenFAST to simulate wind turbines within the turbulent ABL over a selected 10-minute time window, as described in Section 2.2. This second phase takes significantly less simulation time (10 minutes to 1 hour) but requires a finer grid spacing and a smaller timestep to accurately capture small-scale structures and turbine blade displacements. As a result, the turbine simulations have substantially longer walltimes per step compared to the precursor simulation.

In both simulations, the numerical domain (Figure 3.3) extends $3 \text{ km} \times 3 \text{ km} \times 1 \text{ km}$ in the longitudinal (x), lateral (y), and vertical (z) directions, respectively. The refined mesh used for the turbine simulations is initialized with the flow field from the precursor simulation at the starting timestep of the chosen time window. The inflow boundary condition is generated from planes saved during the precursor run, while an outflow boundary condition is applied at the opposite end. To ensure that the turbine models are correctly initialized and the wake is fully developed downstream, an additional 1 minute is simulated before the selected 10-minute window. However, only the 10-minute interval is examined and compared with experimental data during the post-processing phase.

It is important to note that, even though both WTGa1 and WTGa2 are simulated, only data from the leading WTGa1 turbine are processed, as experimental data are available solely for this turbine.

3.2.1 Precursor simulation

The precursor simulation is run with periodic side and inflow/outflow boundary conditions. On the ground, a heat flux equal to 0 K m/s is set, representing the absence of stratification, and a roughness height of 0.01 m is found to effectively mimic real terrain conditions of the SWiFT facility [47]. A constant potential temperature of 304.6 K is applied throughout the domain, except at the upper boundary, where an impermeable free-slip condition with a temperature gradient of 0.003 K/m is imposed to match the gradient in the capping inversion, which lies outside the domain.




Figure 3.3: Precursor simulation domain. The red arrow points to the North, while the black and blue one indicates the streamwise direction.

The simulation is conducted for 20,000 seconds with a timestep of 0.5 seconds to

eliminate initial transients and allow turbulence to fully develop. To accelerate this process, a small perturbation near the ground ($z = 50$ m) is applied to the initial uniform velocity field of $[8.69, 0.0, 0.0]$ m/s, using the method described in [47].

The pressure gradient across the domain is automatically adjusted to create a velocity profile that matches the ensemble average of measured mean velocity at hub-height (in Table 3.1). This is accomplished through the AMR-Wind ABL forcing scheme, which introduces a source term in the momentum equation to drive the flow toward a target velocity at a specific height. Unlike geostrophic forcing, this method allows the desired state to be reached more quickly, providing better control over turbine hub-height conditions, though it may be less effective at higher altitudes [47]. The Coriolis forcing term is computed with the actual latitude of the SWiFT site.

A uniform and isotropic grid resolution of 10 meters is adopted, for a total of 9 million cells. This spacing has already been demonstrated to be appropriate for ABL simulations of this benchmark [47, 50, 51].

3.2.2 Turbine simulations

For the turbine simulations, three different cases (shown in Table 3.2) are carried out, similar to those presented in [50].

Two mesh resolutions, referred to as *coarse* and *fine*, are used. Both use the precursor mesh as background grid. For the coarse mesh, there is a refinement region that halves the spacing from 10 m to 5 m and an inner level where the spacing is 2.5 m (Figure 3.4). For the fine mesh, an additional refinement layer decreases the spacing to 1.25 m (Figure 3.5). The first refinement volume spans 540 m x 140 m x 100 m, starting 190 m in front of the first turbine and at the ground, with each subsequent internal level being 10 m smaller on all sides except for the bottom surface.

All simulations run for 660 seconds, with the timestep (in Table 3.2) carefully selected after a trade-off to ensure that the Courant–Friedrichs–Lewy (CFL) condition is satisfied while also keeping the computational cost acceptable. However, for the ALM simulation, the timestep is far more restricted by the blade travel condition, which requires that the blade tip displacement within a single timestep not exceed one grid spacing [52].

The OpenFAST model of WTGa1 and WTGa2 is provided under the folder `wt_data/v3/OpenFAST_model` in the benchmark GitHub repository [53]. The model was developed by C. Kelley and J. White from Sandia National Laboratories [54] and progressively tuned to match the measurements data. It includes all the aerodynamic, structural and control data required by OpenFAST to compute turbine metrics such as generator power, torque, and rotor speed. Both modified V27s operate in below-rated conditions under variable-speed control, with yaw and

Table 3.2: Description of turbine simulations. All cases are run with AMR-Wind/OpenFAST, as detailed in Section 2.2.

Case	Min. grid spacing [m]	Num. of grid elements	Sim. timestep [s]
ADM coarse	2.5	9.4e6	0.2
ADM fine	1.25	11.5e6	0.1
ALM fine	1.25	11.5e6	0.02




Figure 3.4: Refinement levels for *ADM coarse* turbine simulation.

pitch control turned off. The yaw heading is 0° , though minor deviations from zero in the mean yaw offset occur due to aleatoric uncertainty. The collective pitch values are -0.75° and -2.75° for the upwind and downwind turbines, respectively.

In these simulations we employ the Filtering Lifting Line Correction (FLLC) [55] to account for a non-optimal value of smoothing length scale ϵ introduced in Equation 2.16. Indeed, for the SWiFT Vestas V27 turbines an optimal value of ϵ




Figure 3.5: Refinement levels for *ADM fine* and *ALM fine* turbine simulations.

would lead to extra-fine meshes, resulting in unpredictable computational times [56].

Table 3.3 shows the cost estimates for all the considered cases. We introduce the cost as a measure of computational efficiency, since the simulations were performed with varying numbers of cores based on their availability on the machine at the time, making walltime an unfair metric.

Table 3.3: Cost estimates for turbine simulations. The cost is calculated by multiplying the computation hours by the number of cores allocated, according to [50].

Case	Walltime [hours]	Cores	Cost [Corehours]
ADM coarse	2.3	96	220.8
ADM fine	8.8	64	561.7
ALM fine	26.5	96	2545.8

3.3 Results

The accuracy of the numerical simulations is assessed using experimental inflow, turbine and wake data. To reduce computation time, the validation of inflow and turbine numerical results is done by comparing only to ensemble-averaged data (Table 3.1). The wake analysis is conducted on a fixed reference frame (x , y , z) centered at the turbine hub, with x aligned with the mean wind direction, z pointing upwards, and y completing the right-handed coordinate system.

3.3.1 Atmospheric inflow

The convergence of the precursor simulation is assessed by tracking the evolution of two flow statistics: the horizontal plane-averaged velocity $\langle U_{horiz} \rangle$ and the plane-averaged resolved turbulence intensity $\langle TI_R \rangle$. U_{horiz} and TI_R are defined as follows:

$$U_{horiz} = \sqrt{u^2 + v^2}, \quad (3.2)$$

$$TI_R \equiv \frac{\sqrt{\frac{2}{3} \langle k_{res} \rangle}}{|U|}, \quad (3.3)$$

where u and v are the horizontal wind speed components, $|U|$ is the wind speed magnitude and k_{res} is the resolved-scale TKE.

Figure 3.6 shows the time histories of $\langle U_{horiz} \rangle$ at different altitudes in the domain. As depicted, this metric converges very quickly near the ground, achieving a quasi-steady-state solution in less than 2500 seconds, whereas it takes significantly longer near the top of the boundary layer. There are two primary reasons for this behavior. First, we enforce the ABL velocity profile to match the experimental value at hub height, which is relatively close to the terrain. In contrast, we lack measurements in the upper region of the ABL. Consequently, faster convergence and potentially more accurate results are observed near the surface, as anticipated in Section 3.2.1. Second, to accelerate turbulence development, we initialize the ABL with a small perturbation near the ground. This perturbation is progressively damped and becomes less effective towards the upper boundary of the domain, as shown in Figure 3.7. In this work, our primary focus is on validating the simulation results against data from the SWiFT facility. Therefore, statistical convergence of the entire ABL is not essential. What is critical, however, is the convergence of mean and higher-order statistics within the turbine region, which we consider to extend up to approximately 100 m in altitude.

Figure 3.8 and Figure 3.9 present $\langle U_{horiz} \rangle$ and $\langle TI_R \rangle$ vertical profiles. Figure 3.8 confirms that $\langle U_{horiz} \rangle$ reaches a quasi-steady-state solution relatively quickly. In contrast, Figure 3.9 demonstrates that $\langle TI_R \rangle$ takes considerably longer to converge and has not fully converged even by the end of the precursor simulation. However,




Figure 3.6: Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the precursor simulation.




Figure 3.7: Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.

similar to what was observed by Hsieh et al. (2021) [47], we found that the difference in $\langle TI_R \rangle$ values in the turbines region at $t = 16,000\text{ s}$ and $t = 19,000\text{ s}$ is negligible




Figure 3.8: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.




Figure 3.9: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.

in the context of this validation. Therefore, we conclude that the time window between $t = 16,000$ s and $t = 20,000$ s is appropriate for calculating ABL statistics.

The time-averaged vertical profiles of $\langle U_{horiz} \rangle$ and $\langle TIR \rangle$ over the aforementioned time window are depicted in Figure 3.10 and Figure 3.11, respectively. The simulated mean vertical wind shear profile closely matches the one obtained with Nalu-Wind by Hsieh et al. (2021) [47] and the experimental data from the five sonic anemometers mounted on METa1. However, the simulation slightly underpredicts the mean wind speed below the hub height and overpredicts it above. This is further evidenced by the fact that the wind shear exponent, calculated according to Equation 3.1, is approximately 0.17. Although this differs slightly from the experimental ensemble average of α , it closely matches the values observed in two of the six experimental bins, making it an acceptable estimate in this context (Table 3.1). Finally, the $\langle TIR \rangle$ profile also aligns well with the findings in [47], though it slightly overestimates the experimental data at hub height.




Figure 3.10: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [16,000; 20,000]$ s. The orange line is reconstructed from the study by Hsieh et al. (2021) [47]. The black squares are the measurements collected by METa1 meteorological tower [53]. The dashed black line indicates the rotor hub-height.




Figure 3.11: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, averaged over the time window $t = [16,000; 20,000]$ s. The orange line is reconstructed from the study by Hsieh et al. (2021) [47]. The dashed black line indicates the rotor hub-height.





Figure 3.12: Time series of hub-height wind speed, U_{HH} . The blue line represents the output from AMR-Wind, while the gray line is an overlay of the six individual experimental time series recorded by METa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.

3.3.2 Wind turbine response


Herein, we compare three quantities of interest for the WTGa1 turbine between the experimental data and the simulation results: rotor speed, generator power, and generator torque.

Figure 3.13, Figure 3.14 and Figure 3.15 display the time series of WTGa1 generator power and rotor speed obtained from the three simulation strategies in AMR-Wind/OpenFAST, compared with both the six experimental time series and the results from [47] using Nalu-Wind/OpenFAST. A 30-second moving average is applied to all time series. As can be noted, AMR-Wind/OpenFAST better captures the end of the experimental rotor speed and generator power time series compared to Nalu-Wind/OpenFAST. This discrepancy is likely due to Nalu-Wind showing a decrease in the hub-height wind speed, U_{HH} , during that period (Figure 3.12), which does not align with the experimental data. In contrast, AMR-Wind displays a similar behavior in the interval [100, 200] s. This discrepancy must be attributed only marginally to differences between the two codes, and is primarily due to the different 10-minute time windows selected in the two studies for turbine simulations.

The mean and standard deviation of the simulated and experimental time series are detailed in Table 3.4 and represented in the bar chart of Figure 3.16. All the simulation cases have very similar mean rotor speeds, while there is more variation in the generator power and torque across the different approaches, as also noted by Doubrawa et al. (2020) [51]. Nevertheless, all simulated quantities remain within one standard deviation of the experimental data. Compared to the measurements, the simulated WTGa1 generator power and torque tend to overestimate, as also observed by Hsieh et al. (2021) [47]. This is particularly pronounced in the ALM case, as also found by Hsieh et al. (2020) [50]. One potential reason is that both AMR-Wind/OpenFAST and Nalu-Wind/OpenFAST time series of simulated rotor speed fail to capture some significant drops in the experimental rotor speed, resulting in a mean simulated rotor speed higher than the ensemble average of the experimental 10-minute mean rotor speed. This idea is supported by the fact that the *ALM fine* case exhibits the smallest drop in the rotor speed and also the highest power (Figure 3.15), suggesting that the controller might play a significant role in the overestimation of the simulated generator power. Other possible causes could include uncertainties in experimental measurements, poor or inadequate mesh resolution, inaccuracies in the OpenFAST turbine model, or errors in the coupling process between AMR-Wind and OpenFAST.




(a)




(b)

Figure 3.13: Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with *ADM coarse* approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.




(a)




(b)

Figure 3.14: Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with *ADM fine* approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.



(a)



(b)

Figure 3.15: Time series of generator power (a) and rotor speed (b). The blue line represents the output from AMR-Wind with *ALM fine* approach, while the gray line is an overlay of the six individual experimental time series recorded by WTGa1 [53]. The orange line shows the reconstructed time series from [47]. The dashed black line represents the experimental mean value, while the gray shaded region indicates the range of one standard deviation around it.

Table 3.4: Mean \pm standard deviation for time series of WTGal response quantities of interest.

	Rotor Speed [RPM]	Gen. Power [kW]	Gen. Torque [Nm]
ADM coarse	43.45 ± 0.34	84.8 ± 12.7	701.0 ± 102.3
ADM fine	43.44 ± 0.39	85.1 ± 13.7	703.4 ± 110.7
ALM fine	43.54 ± 0.15	93.2 ± 14.8	769.5 ± 120.7
Experiment	42.76 ± 1.70	79.1 ± 20.1	622.5 ± 144.7




Figure 3.16: Mean (colored bar) and standard deviation (horizontal black line) for time series of wind turbine response quantities.

3.3.3 Wind turbine wake

The wake analysis aims to evaluate the ability of each simulation approach to accurately reproduce wake strength, recovery, shape and expansion. We investigate the fields of normalized velocity deficit, calculated as

$$U_d(t, x, y, z) = \frac{U(t, x, y, z)}{U_\infty} - 1, \quad (3.4)$$

and resolved turbulence intensity in the rotor near-wake, usually defined as “the region immediately downwind of the turbine with a length of 2–4 rotor diameters” [57]. U represents the horizontal wind speed and U_∞ denotes the time-averaged freestream wind speed profile sampled at $x/D = -2.5$, where METa1 is located.




Figure 3.17: Time-averaged normalized velocity deficit in YZ-planes located at $x/D = 2, 3, 4$ and extending 1D from the hub center in both positive and negative lateral and vertical direction. Results are shown for *ADM coarse*, *ADM coarse* and *ADM coarse* cases. The black circle outlines rotor area.




Figure 3.18: Resolved turbulence intensity in YZ-planes located at $x/D = 2, 3, 4$ and extending 1D from the hub center in both positive and negative lateral and vertical direction. Results are shown for *ADM coarse*, *ADM coarse* and *ADM coarse* cases. The black circle outlines rotor area.

Figure 3.17 and Figure 3.18 show the time-averaged normalized velocity deficit and resolved turbulence intensity in the turbine wake for each of the three simulated cases. The typical phenomenology of wind turbine wake is evident. Immediately behind the rotor there is a strong reduction in wind speed, which progressively recovers with increasing downwind distance. Concurrently, the wake expands in the cross-sectional planes due to the entrainment of the outer flow, as reflected by the spread of velocity deficit and resolved turbulence intensity. As the wake widens, its shape also changes with the downstream position. The classical ring pattern, distinctly visible at $x/D = 2$, gradually smears downwind.

To conduct a more quantitative analysis, we represent the time-averaged normalized velocity deficit and resolved turbulence intensity along the two symmetry

lines of the YZ-planes (Figure 3.19 and Figure 3.20).




Figure 3.19: Horizontal profiles of normalized velocity deficit (a) and resolved turbulence intensity (b) averaged in time and sampled at hub height for different downstream distances. The black dashed lines indicate the rotor edges.




Figure 3.20: Vertical profiles of normalized velocity deficit (a) and resolved turbulence intensity (b) averaged in time and sampled at hub height for different downstream distances. The black dashed lines indicate the rotor edges.

As expected, resolved turbulence intensity peaks near rotor edges, where mean shear and turbulent fluxes are intense. These regions roughly correspond to the side shear layers, where the wake flow interacts and mixes with the surrounding freestream flow, as clearly shown by the velocity deficit profiles. Then, the side shear layers expand further downstream and the velocity profile gradually recovers and becomes more uniform. In addition, from Figure 3.19b and Figure 3.20b it is evident that the resolved turbulence intensity increases downwind in the near-wake, as well-known in the literature [57].




Figure 3.21: Maximum of time-averaged velocity deficit fields as a function of the downstream position. PALM, Nalu-Wind and SpinnerLidar data are taken from [51].

The horizontal and vertical profiles of the velocity deficit are in good agreement among the three AMR-Wind/OpenFAST models, although a larger discrepancy can be observed around the wake center at $x/D = 2$, especially between the two actuator models. Notably, the *ALM fine* case exhibits a greater velocity deficit in this region, according with the overestimation of generator power and torque observed in Section 3.3.2. On the contrary, resolved turbulence intensity has more variability along the lateral and vertical directions, particularly within one rotor area ($-0.5D < y, z < 0.5D$). By comparing our results with the scanning lidar measurements [51], we find that the three simulation approaches using AMR-Wind/OpenFAST accurately capture the wake strength, except at $x/D = 2$, where

it is underestimated. However, this is not surprising, as such models tend to struggle in modeling the region right behind the turbine. Moreover, for this benchmark, other studies employing LES solvers have reported a similar underestimation of the velocity deficit [47, 51], making this result acceptable within the context of the comparison. Conversely, the wake recovery rate appears to be underestimated not only compared to experimental measurements but also relative to other LES solvers, such as PALM and Nalu-Wind. The wake recovery is significantly influenced by both the turbulence in the incoming flow and shear-generated turbulence. While the former is accurately captured in the precursor simulation, the latter presents much more uncertainties (e.g spatial and temporal resolution, the AMR-Wind actuator model and its coupling algorithm with the OpenFAST turbine model, etc) and its underprediction might weaken the wake recovery. Hence, further investigation into these and other potentially related aspects is recommended. Figure 3.21 provides a clear representation of these trends.

Chapter 4

Impact of floating solar PV on wind farm flow and power output

The presence of FSPV panels within offshore wind farms introduces an obstacle to the flow reaching the turbines. Intuitively, this might suggest a negative impact on the wind farm operations due to the drag imposed by the obstacle. Nevertheless, recent studies have shown that, under certain specific conditions, the presence of obstacles or rougher surfaces can increase on- and off-shore wind farm power generation due to the enhanced turbulent mixing, which speeds up wake recovery and improves downstream turbine performance.

Shamsoddin and Porté-Agel (2017) [58], Liu and Stevens (2020) [59] used LES and IBM [60, 61] to analyze the performance of wind farms situated on hilly terrain. The results from [58] indicated that LES can accurately replicate the flow field, particularly the speedup over the hilltop, as well as the velocity deficit and the turbulence intensity enhancement induced by the turbines. Furthermore, Liu and Stevens (2020) [59] highlighted that turbines significantly taller than the hill generally benefited from the speedup over the hill, while the performance of shorter turbines was highly dependent on their position relative to the hill. Power generation tended to peak at the hilltop, with the lowest output typically occurring just downstream of the hill.

Ge et al. (2021) [62] explored the performance of two wind turbines positioned behind a building-like, wall-attached cube using LES and IBM. Their study revealed that, for specific distances between the first turbine and the cube, the total power output of the two turbines surpassed that of the setup without the cube, despite the first turbine experienced a significant power loss. This was due to the second turbine benefiting from a faster wake recovery of the first turbine, facilitated by

the enhanced turbulent transport and secondary mean flow structures induced by the cube.

Other research employed surface roughness heterogeneities to account for the spatial variability of the momentum fluxes imposed by the ground surface on the ABL and wind farm flow. Kethavath and Ghaisas (2024) [63] conducted LES of the flow over a wind farm sited downstream an abrupt rough-to-smooth surface roughness jump. They observed that wind farms located behind the surface roughness jump generated more power compared to those placed on homogeneously rough surface. This increase was mainly because of the higher power output of the downstream turbines and only slightly due to the increased power of the first-row turbine.

Tobin and Chamorro (2017) [64], Liu and Stevens (2021) [65] used an increase in surface roughness to model windbreaks in LES simulation of large wind farms. Tobin and Chamorro (2017) [64] demonstrated that windbreaks had a negative effect on the performance of an infinite wind farm. Indeed, the additional drag introduced by the windbreaks outweighed the benefits of the speedup over them, ultimately reducing the wind speed reaching the turbines. Conversely, Liu and Stevens (2021) [65] implemented windbreaks in a finite wind farm and observed that, in specific configurations, they actually enhanced overall power production of smaller wind farms.

In this chapter, we utilize the IBM integrated into AMR-Wind to represent FSPV panels within a benchmark offshore wind farm. We assess the panels' impact on the wind farm's performance, focusing on both flow dynamics and power output. Specifically, building on the works [59, 62, 63, 64, 65], we explore whether the installation of FSPV panels can enhance wind farm power extraction.

4.1 Benchmark wind farm simulation

The panels' impact is assessed in neutral stratification for a reference aligned wind farm consisting of 4 rows and 4 columns in x and y directions, respectively (Figure 4.1). The coordinate system is identical to that used in Chapter 3. The turbine model used in this study is the IEA Wind 15-MW offshore reference wind turbine [66], which has a rotor diameter of $D = 240$ m and a hub height of $z_h = 150$ m (Figure 4.2). The 16 wind turbines are implemented within a domain measuring $9,920$ m \times $6,720$ m \times $1,040$ m in the longitudinal, lateral, and vertical directions, respectively. The first row of wind turbines is located approximately $8D$ downstream of the inflow, and both the streamwise and spanwise turbine spacings, S_x and S_y , are set to $7D$, which is typical of that used in wind farms and in previous similar studies [67, 68].




Figure 4.1: Benchmark wind farm layout.

4.1.1 Simulation setup

For this analysis, we adopt the same simulation strategy detailed and validated in Chapter 3. Any information omitted in this section is assumed to remain unchanged.

The simulation workflow is divided into a precursor and a subsequent wind farm simulation. The precursor simulation runs for 6 hours with a timestep of 0.8 seconds, and the initial velocity field is uniformly set to $[8.0, 0.0, 0.0]$ m/s. The same velocity vector is also imposed at hub height throughout the entire simulation. The sea surface is modeled using an aerodynamic roughness height $z_0 = 0.001$ m [38, 69]. A constant potential temperature of 285 K is enforced across the entire domain, which represents a reasonable value for offshore wind sites. The Coriolis source term is calculated based on the mean latitude of the wind farms within the first Belgian offshore wind energy zone. The computational domain is discretized into a uniform, isotropic grid of $496 \times 336 \times 52$ cells. This level of grid resolution have been shown to provide reliable results in previous similar works [62, 63, 65].

The wind farm simulation is conducted in the same domain, adding an extra resolution level that uniformly halves the spacing from 20 m to 10 m (Figure 4.3).




Figure 4.2: IEA Wind 15-MW offshore reference wind turbine. Figure reprinted from Ref. [66].

The simulation runs for 40 minutes, which is considered sufficient time for the convergence of the flow and turbine statistics of interest, and advances with a timestep of 0.4 s to ensure that the CFL condition is met. In AMR-Wind, the turbine is modeled using only ADM. This choice is driven by two factors: first, ADM provides a better power prediction than ALM for relatively coarse meshes, like those we use in this analysis (see Section 3.3.2); second, simulating an entire wind farm with ALM would require spatial and temporal resolutions that are beyond our computational resources. The OpenFAST model of the wind turbine is retrieved from the dedicated GitHub repository [70]. The turbines operate under variable-speed control and with yaw and pitch controls disabled, experiencing

below-rated wind speeds due to the inflow conditions generated by the precursor simulation (Figure 4.4).




Figure 4.3: Refinement level for wind farm simulation.




Figure 4.4: Turbine controller regulation trajectory. Figure adapted from Ref. [66].

4.1.2 Precursor simulation results

As illustrated in Chapter 3, we examine the time histories of the horizontally plane-averaged velocity, $\langle U_{horiz} \rangle$, and the plane-averaged resolved turbulence intensity, $\langle TI_R \rangle$, to assess the convergence of the precursor simulation.




Figure 4.5: Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the precursor simulation.




Figure 4.6: Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.




Figure 4.7: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.




Figure 4.8: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the precursor simulation. The dashed black line indicates the rotor hub-height.

Figure 4.5, Figure 4.6, Figure 4.7 and Figure 4.8 demonstrate that, after approximately $t = 19,200$ s, both $\langle U_{horiz} \rangle$ and $\langle TI_R \rangle$ reach a quasi-steady-state solution across nearly the entire domain, particularly within the wind farm's operational area. Therefore, the ABL statistics are computed over the time window between t

= 19,200 s and $t = 21,600$ s, corresponding to the last 40 minutes of simulation. The same time window is used to generate the inflow for the subsequent wind farm simulation. Figure 4.9 and Figure 4.10 show the time-averaged vertical profiles of $\langle U_{horiz} \rangle$ and $\langle TIR \rangle$.

For further details on precursor simulation post-processing, the reader is referred to subsection 3.3.1.




Figure 4.9: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [19,200; 21,600]$ s. The dashed black line indicates the rotor hub-height.




Figure 4.10: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TIR \rangle$, averaged over the time window $t = [19,200; 21,600]$ s. The dashed black line indicates the rotor hub-height.

4.1.3 Wind farm simulation results

With finite-size simulations one studies the so-called *wind farm flow-development region* [57], which approximately starts from the first row of turbines. Downwind of the wind farm's leading edge, momentum extraction by the wind turbines leads to the formation of turbine-wake flows that support the development of an IBL, which grows with downwind distance. Downstream of the wind farm's trailing edge, the wake flow of the wind farm gradually recovers momentum, accelerating in the streamwise direction. This phenomenology is clearly illustrated in Figure 4.11 and Figure 4.12, which depict the streamwise velocity field at the end of the simulation.




Figure 4.11: Streamwise velocity field in the turbine hub-height plane at $t = 21,600$ s. The dashed black lines divide the plane into four equal bands across which the spatial average is computed.




Figure 4.12: Streamwise velocity field in the symmetry XZ-plane of the third column turbines at $t = 21,600$ s.

The evolution of the IBL due to the presence of the wind farm is also evident in Figure 4.13, Figure 4.14, Figure 4.15 and Figure 4.16. Within the IBL, there is a

significant reduction in wind speed and a net increase in turbulent kinetic energy with respect to the ABL values in the absence of the wind farm (corresponding to $t = 19,200.4$ s). This is because of the farm's blockage effect, which also induces a speedup over the wind farm (above $z \sim 500$ m). Since the figures represent averages over the horizontal plane, these effects may be even more pronounced locally.




Figure 4.13: Time histories of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several altitudes during the wind farm simulation.




Figure 4.14: Time histories of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several altitudes during the precursor simulation.




Figure 4.15: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, at several times during the wind farm simulation. The dashed black line indicates the rotor hub-height.




Figure 4.16: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, at several times during the wind farm simulation. The dashed black line indicates the rotor hub-height.

Approximately after $t = 20,400$ s, both $\langle U_{horiz} \rangle$ and $\langle TI_R \rangle$ reach a quasi-steady state, thus we select the last 20 minutes of the simulation ($t = [20,400; 21,600]$) to compute the averages profiles (Figure 4.17 and Figure 4.18). The same time window is used for following statistical analysis of the flow.




Figure 4.17: Vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, averaged over the time window $t = [20,400; 21,600]$ s. The dashed black line indicates the rotor hub-height.




Figure 4.18: Vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, averaged over the time window $t = [20,400; 21,600]$ s. The dashed black line indicates the rotor hub-height.

The flow field is examined looking at the time-averaged streamwise velocity deficit and the most considered turbulence quantities in the literature [57], i.e. the TKE, streamwise turbulence intensity and turbulent momentum fluxes in the lateral and vertical directions $\overline{\rho u'v'}$ and $\overline{\rho u'w'}$, respectively. The velocity deficit is normalized by the mean incoming wind speed at hub level, \bar{U}_h , while the turbulence statistics by \bar{U}_h^2 .

We represent the grid-resolved part of the investigated quantities in XY and XZ planes, averaging along the y -direction due to the periodic conditions set at the two lateral boundaries. The fields in the sample XY-plane result from averaging across four $7D$ -wide strips that divide the hub-height plane (Figure 4.11), while the ones in the sample XZ-plane are derived from averaging across the midplanes (vertical symmetry planes) of the four turbine columns.

Figure 4.19 and Figure 4.20 show the normalized streamwise velocity deficit $(U_{ABL} - U)/\bar{U}_h$, where U_{ABL} is the incoming wind speed profile. Immediately upstream of the wind farm, a region known as *induction zone* [57] can be recognized, where the approaching boundary-layer flow decelerates and deflects vertically and laterally due to the blockage effect caused by the wind farm as a whole. This zone is followed by the flow-development and wake regions described earlier.




Figure 4.19: Normalized streamwise velocity deficit averaged over time and across the $7D$ -wide XY-strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.




Figure 4.20: Normalized streamwise velocity deficit averaged over time and across the symmetry XZ-planes of the turbine columns (indicated by solid black lines).

Figure 4.21 and Figure 4.22 illustrate the resolved TKE field. The TKE is generally high within the wind farm due to turbine wake-flows, especially near the rotor edges, where wind shear and turbulent fluxes are stronger. The maximum of TKE occurs at the rotor top-tip, while it is suppressed near the ground. Moreover, in the XY-plane, lateral asymmetry in TKE production is observed, caused by the wake rotation which is opposite to the rotor rotation. Similar comments can be made for the streamwise turbulence intensity (Figure 4.23 and Figure 4.24).




Figure 4.21: Normalized turbulent kinetic energy averaged across the $7D$ -wide XY -strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.




Figure 4.22: Normalized turbulent kinetic energy averaged across the symmetry XZ -planes of the turbine columns (indicated by solid black lines).




Figure 4.23: Normalized streamwise turbulence intensity averaged across the $7D$ -wide XY -strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.




Figure 4.24: Normalized streamwise turbulence intensity averaged across the symmetry XZ -planes of the turbine columns (indicated by solid black lines).

The spatial distribution of the Reynolds shear stresses (Figure 4.25 and Figure 4.26), $-\overline{u'v'}$ and $-\overline{u'w'}$, contributes to the entrainment of kinetic energy from the outer flow toward the wake center, which supports the IBL expansion. Similar to the TKE and streamwise turbulence intensity, the magnitude of the turbulent fluxes is greater at the rotor tips, particularly close to the upper edge of the wake where the wind shear is more intense. These momentum fluxes increase downstream, resulting in enhanced wake recovery for subsequent turbine rows, as evidenced by the streamwise velocity deficit field (Figure 4.19 and Figure 4.20).




Figure 4.25: Normalized kinematic lateral turbulent momentum flux averaged across the $7D$ -wide XY -strips dividing the hub-height plane into four sections. The solid black lines represent the turbines.




Figure 4.26: Normalized kinematic vertical turbulent momentum flux averaged across the symmetry XZ -planes of the turbine columns (indicated by solid black lines).

4.2 Sensitivity analysis on floating solar PV layouts

Here, we conduct a sensitivity study on the impact of different FSPV layouts on wind farm performance, similar to the approach adopted by Liu and Stevens (2021) [65] with windbreaks. Table 4.1 summarizes all the cases considered in this analysis. Case 0 indicates the farm-only scenario discussed in Section 4.1.3, while Case 1 uses a baseline FSPV arrangement, from which the other case studies are built.

Table 4.1: Overview of the cases performed within the sensitivity study. Here, h , l and w are the height, length and width of the FSPV platform, respectively. The distance between the platform center and the wind turbine is denoted by x_c , and N_{PV} indicates the number of platforms. P/P_0 is the ratio of the wind farm power output with FSPV panels (P) to that of the wind farm without FSPV panels (P_0).

Case	h/z_h	l/D	w/D	x_c/D	N_{PV}	P/P_0
0	-	-	-	-	0	1
1	0.13	5	1	3.5	16	0.87
2	0.26	5	1	3.5	16	0.75
3	0.065	5	1	3.5	16	0.91
4	0.13	6.5	1	3.5	16	0.85
5	0.13	3.5	1	3.5	16	0.91
6	0.13	1	5	3.5	16	0.89
7	0.13	1	5	1.5	16	0.91
8	0.13	1	5	5.5	16	0.87
9	0.13	5	1	3.5	4	0.98

4.2.1 Power budget analysis

The general and most evident result of this analysis is that, in all cases, the presence of the panels leads to an overall reduction in wind farm power extraction. From Table 4.1, it becomes clear that this reduction is more important as the number of panels increases or as they become larger, while reducing the distance between panels and wind turbines is beneficial for power production.

In addition to assessing the overall impact of FSPV panels on wind farm production, we are also interested in understanding their influence on each individual row. Figure 4.27 shows the time-averaged normalized power production P/P_0 as a function of wind farm row for all cases. Here, P indicates the power averaged over all columns of a given row, while P_0 denotes the power averaged across all columns of the first row in the farm-only scenario. We normalize by P_0 as it reflects the mean power output of a single row under the most undisturbed case. The typical knee in the P/P_0 vs. row curve can be observed [65, 71]. Beyond the overall reduction in wind farm power output due to the presence of FSPV panels (Table 4.1), this plot shows that each individual row also experiences a decrease in power.




Figure 4.27: Normalized power output P/P_0 as a function of the wind farm rows for all cases analyzed within the sensitivity study.

In Figure 4.28 and Figure 4.29, we isolate the effect of FSPV platforms height and length, respectively. It can be concluded that the parameter h/z_h has a much greater influence than l/D , which is expected given the significant role of vertical flow motion in the ABL.




Figure 4.28: Normalized power output P/P_0 as a function of the wind farm rows for different values of the h/z_h parameter.




Figure 4.29: Normalized power output P/P_0 as a function of the wind farm rows for different values of the l/D parameter.

Cases 6, 7 and 8 are obtained from Case 1 by swapping the length and width dimensions of the FSPV platforms. Figure 4.30 shows that positioning the panels closer to the wind turbines enhances the power production of each row. This finding suggests a speedup effect over the panels, which can be exploited to increase turbine power output.




Figure 4.30: Normalized power output P/P_0 as a function of the wind farm rows for different values of the x_c/D parameter.

4.2.2 Flow analysis

To understand the reasons behind the numerical results of Section 4.2.1, we need to examine how FSPV panels alter the flow conditions in which the turbines operate.

On one hand, the panels induce a reduction in wind speed (Figure 4.31), with the most significant decrease occurring around the rotor area. Since power generation is proportional to the cube of the wind speed ($P \propto U^3$), this reduction is the primary cause of power loss.




Figure 4.31: Time-averaged vertical profiles of horizontal plane-averaged ABL velocity, $\langle U_{horiz} \rangle$, for all cases. The dashed black line indicates the rotor hub-height.

On the other hand, the turbulent kinetic energy increases overall within the wind farm, especially near the lower part of the domain (Figure 4.32). This increase in turbulence is not necessarily deleterious to power generation; in fact, it is well-known that higher atmospheric turbulence intensity can enhance wake recovery [57]. However, a higher turbulence level can have a detrimental effect on the lifespan of the blades due to increased fatigue, though we do not take this into account in the present study.

For Case 2 (Figure 4.33) and Case 7 (Figure 4.34) we also analyze the flow fields, subtracting the corresponding fields in the only-farm scenario. This helps us to visualize how the addition of the panels affects the flow within and around the wind farm. We choose Case 2 because it exhibits the most significant reduction in power production due to the presence of the panels, which should make the associated phenomena more evident. Additionally, Case 7 is analyzed to investigate the reasons behind the increase in power output with reduced spacing between the panels and the turbines.

We first examine Case 2. Figure 4.35 and Figure 4.36 clearly illustrate the

increase in velocity deficit caused by the panels, which intensifies downstream.




Figure 4.32: Time-averaged vertical profiles of plane-averaged ABL resolved turbulence intensity, $\langle TI_R \rangle$, for all cases. The dashed black line indicates the rotor hub-height.




Figure 4.33: Side view (top) and top view (bottom) of FSPV layout in Case 2.




Figure 4.34: Side view (top) and top view (bottom) of FSPV layout in Case 7.




Figure 4.35: Delta of normalized streamwise velocity deficit in the hub-height XY-plane, obtained by subtracting the Case 0 field from the Case 2 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case.

The panels generate a region of speedup located at the front, where the flow accelerates and deflects upward, while a boundary layer develops over the panels. The Case 2 layout is particularly disadvantageous for power generation as the turbines operate in the wake of the panels rather than in the speedup region, which is located right behind each rotor in this case (Figure 4.35, Figure 4.36 and Figure 4.37). Additionally, the vertical entrainment of the external flow toward the center of the wake is diminished in the presence of the panels (Figure 4.38), leading to a slower wake recovery.




Figure 4.36: Delta of normalized streamwise velocity deficit in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case. Light brown areas indicate out-of-scale values.




Figure 4.37: Delta of normalized turbulent kinetic energy in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Light brown areas indicate out-of-scale values.




Figure 4.38: Delta of normalized kinematic vertical turbulent momentum flux in the turbine midplane, obtained by subtracting the Case 0 field from the Case 2 field. Light brown areas indicate out-of-scale values.

In contrast, with the layout of Case 7 the turbines operate in proximity of the speedup region generated by the panels (Figure 4.39, Figure 4.40 and Figure 4.41). This explains why the power output increases when the distance between the panel and turbine decreases. However, vertical entrainment of kinetic energy is still reduced compared to Case 0 (Figure 4.42).




Figure 4.39: Delta of normalized streamwise velocity deficit in the hub-height XY-plane, obtained by subtracting the Case 0 field from the Case 7 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case.




Figure 4.40: Delta of normalized streamwise velocity deficit in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field. Positive values indicate lower velocities induced by the panels with respect to the only-farm case. Light brown areas indicate out-of-scale values.




Figure 4.41: Delta of normalized turbulent kinetic energy in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field.




Figure 4.42: Delta of normalized kinematic vertical turbulent momentum flux in the turbine midplane, obtained by subtracting the Case 0 field from the Case 7 field.

Based on these results, we can conclude that the negative impact of the wake and the weakening in wake recovery generated by FSPV panels outweighs the beneficial effect of the speedup above them, resulting in a flow deceleration and therefore a reduction in the power extracted by the turbines. To mitigate the negative impact of the panels, it is essential to place the turbines closer to the speedup region, as the comparison between Case 2 and Case 7 revealed. Furthermore, a layout similar to that of Case 7 suggests that a positive impact of the panels on power output could even be achieved if designs more similar to windbreaks are adopted, for which [65] have already demonstrated feasibility.

Chapter 5

Conclusions

In this thesis, we experimented with the AMR-Wind/OpenFAST framework to conduct LES of wind farms. Additionally, we assessed the impact that the implementation of FSPV panels within an offshore wind farm might have on the power extraction of the wind turbines.

The strong and weak scaling tests, performed with AMR-Wind on VKI Ealin43 CPUs and applied to a reference ABL simulation for wind farm modeling, both showed that the code scales almost ideally up to approximately 16 cores. Past this limit, the speedup still increases, though with a notable decline in parallelization efficiency.

Ideally, such scaling tests should be conducted a priori for each specific case study to optimize the use of computational resources. In our case, we performed only this analysis primarily due to time constraints; however, there are no great differences between the ABL simulations carried out in this thesis, either in terms of problem size, physics or boundary conditions. Therefore, the results of this scalability test were used to inform our resource allocation decisions for all the ABL simulations conducted in this work.

The investigation of AMR-Wind performance on VKI Ealin43 GPUs could be an attractive topic for future research, given the growing importance of GPU acceleration.

Before simulating large wind farms, the LES methodology employed throughout this study was first validated against other similar LES solvers and experimental measurements within the SWiFT Neutral Benchmark. The tested methodology consists of two stages: a precursor step, where a fully developed turbulent ABL is established, and a subsequent step, where the wind turbines are integrated into the simulation. Two actuator models in AMR-Wind (ADM and ALM) were used for the turbines, coupled with the OpenFAST turbine model, and two grid resolutions (coarse and fine) were tested. The accuracy of the simulations was evaluated in terms of atmospheric inflow, wind turbine response, and near-wake characteristics.

The ABL was captured correctly, particularly in the region where the wind turbine operates, as demonstrated by the strong agreement with experimental and literature results on horizontal wind speed and turbulence intensity.

Regarding the wind turbine response, the rotor speed showed low sensitivity to both the actuator model and the grid resolution, likely due to the type of controller used, and matched well with experimental data as well as previous studies using other solvers. In contrast, generator power and torque exhibited higher variability and a tendency to be overestimated compared to experimental data, especially with the ALM, although they remained within one standard deviation of measurements. One possible explanation is that the simulated rotor speed was slightly higher than the measured value, probably because of an imperfect controller performance. However, other factors may be related to this behavior, such as errors in the coupling between the actuator model in AMR-Wind and the turbine model in OpenFAST, insufficient grid resolution, or uncertainties in the experimental measurements. While our results are consistent with most of the studies using LES within the same benchmark, these factors warrant further investigation in future studies.

The wake analysis of the turbine was conducted both qualitatively, by visualizing the time-averaged normalized velocity deficit and resolved turbulence intensity on YZ planes positioned downstream of the turbine in the near-wake, and quantitatively, by plotting the horizontal profile at hub height and the vertical profile in the turbine midplane for the two aforementioned quantities. The planes in the turbine wake allowed us to observe the typical wake phenomenology downstream of the wind turbine. The profile analysis showed that all the simulated cases are in reasonable agreement for the velocity deficit, although the ALM exhibited a greater reduction in velocity in the core of the wake, consistently with the overprediction of generator power and torque found in the turbine response. As expected and aligned with previous studies, the resolved turbulence intensity displayed more variability, especially within one rotor area. When comparing the results with experimental campaign data and with findings from other LES solvers (e.g. Nalu-Wind, PALM) in similar studies, it was found that both the wake strength and shape were correctly captured, while the wake recovery rate was underestimated. This could be due to an underprediction of the shear-generated turbulence in the turbine wake. However, this hypothesis would need to be confirmed through a more detailed study, as it represents just one of the possible causes.

The validated methodology was then applied to simulate a finite, aligned offshore benchmark wind farm. Subsequently, the FSPV panel platforms were integrated into the simulation domain as block-shaped bodies using the only IBM currently integrated in AMR-Wind. It was not possible to use the actual geometry of the panels, as this IBM is capable of representing only spherical, cylindrical, or block-shaped bodies. However, the research team working on AMR-Wind has two ongoing developments for obstacle representation: an immersed forcing method to

explicitly specify a blockage region through Standard Triangulation Language (STL) or Digital Elevation Model (DEM) files, and a method to impose heterogeneous roughness in the wall model via a user-specified database or a National Land Cover Database (NLCD). Moreover, given these limitations, it was deemed unnecessary to simulate the space between the panels of the same platform, which were thus considered as a single block. For these reasons, it would be interesting to conduct a similar study using the real geometry of the panels.

The sensitivity analysis of the wind farm's performance on the layout of the panels (number, size, orientation, and distance from the turbines) produced the following results. All the analyzed cases reported a power loss for the entire wind farm, which can exceed 10%, as well as for each row of turbines within the farm. This loss became more significant as the number or size of panels increased, because of the drag generated by the panels, which reduces the wind speed reaching the turbines. The resulting velocity deficit worsened with increasing downstream distance. Specifically, the panel height had a greater influence on power compared to the longitudinal and lateral dimensions, as expected, since it alters most the vertical motion of the flow. In contrast, when assessing the effect of orientation and turbine-panel distance, it was observed that power generation improves when the space between the turbine and the panels is reduced, thanks to the speedup above the panels, similarly to what found for windbreaks in previous works. Additionally, for all cases, a decrease in the vertical entrainment of the outer flow towards the core of the turbine wake was also observed, leading to a weakening of the wake recovery. This effect also negatively impacts power generation, as the downstream turbine operates in a more disturbed flow.

In conclusion, this study showed that the negative impact of drag and the weakening of wake recovery caused by FSPV panels outweighs the beneficial effect of the speedup above them, resulting in a deceleration of the flow and, consequently, a reduction in the power extracted by the turbines of the wind farm. However, the number of simulations conducted is relatively small, and increasing the power output of an offshore wind farm by integrating FSPV panels could still be possible, for example, if geometries capable of enhancing the beneficial speedup effect above them were used. Additionally, many variables remain to be considered. Besides the possibility of using a realistic panel geometry or representing their presence through a heterogeneous roughness, it would be valuable to assess the effects of heating caused by the panels, atmospheric stratification, wind direction, wind farm layout (aligned or staggered) or flow regime (finite or infinite wind farm), just to name a few. Finally, since this study identified a significant loss in wind farm power due to the presence of FSPV panels, an energy analysis should also be conducted to evaluate whether, and to what extent, the energy generated by the panels would surpass this loss, thus justifying the investment in integrating the two technologies.

Bibliography

- [1] Intergovernmental Panel on Climate Change (IPCC). *Global Warming of 1.5°C: An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. Ed. by Valérie Masson-Delmotte et al. Cambridge, UK and New York, NY, USA: Cambridge University Press, 2018, p. 616. URL: <https://doi.org/10.1017/9781009157940>.
- [2] International Renewable Energy Agency (IRENA). *World Energy Transitions Outlook 2023: 1.5°C Pathway*. June 22, 2023. URL: <https://www.irena.org/Publications/2023/Jun/World-Energy-Transitions-Outlook-2023>.
- [3] Stefan Schlömer et al. «Annex III: Technology-specific cost and performance parameters». In: *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by Ottmar Edenhofer et al. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2014.
- [4] Global Wind Energy Council (GWEC). *Global Wind Report 2024*. Apr. 15, 2024. URL: <https://gwec.net/global-wind-report-2024/>.
- [5] WindEurope. *Getting fit for 55 and set for 2050*. June 2021. URL: <https://windeurope.org/intelligence-platform/product/getting-fit-for-55-and-set-for-2050/>.
- [6] Global Wind Energy Council (GWEC). *Floating Offshore Wind – a Global Opportunity*. Mar. 2022. URL: <https://gwec.net/wp-content/uploads/2022/03/GWEC-Report-Floating-Offshore-Wind-A-Global-Opportunity.pdf>.
- [7] International Energy Agency (IEA). *World energy outlook 2023*. Licence: CC BY 4.0 (report); CC BY NC SA 4.0 (Annex A). Oct. 24, 2023. URL: <https://www.iea.org/energy-system/renewables/wind>.

- [8] WindEurope. *Wind energy in Europe: 2023 Statistics and the outlook for 2024-2030*. Feb. 2024. URL: <https://windeurope.org/intelligence-platform/product/wind-energy-in-europe-2023-statistics-and-the-outlook-for-2024-2030/>.
- [9] URL: <https://www.belgianoffshoreplatform.be/en/>.
- [10] URL: <https://economie.fgov.be/en/themes/energy/belgian-offshore-wind-energy>.
- [11] International Energy Agency (IEA). *IEA WIND TCP Annual Report 2022: Belgium*. Oct. 2023. URL: https://iea-wind.org/wp-content/uploads/2023/12/Belgium_2022.pdf.
- [12] URL: https://www.elia.be/en/news/press-releases/2024/01/20240129_meshed-high-voltage-grid-in-belgian-north-sea.
- [13] URL: <https://www.offshore-energy.biz/oceans-of-energy-launches-floating-solar-farm-in-belgium/>.
- [14] URL: <https://www.seavolt.be/>.
- [15] URL: <https://oceansofenergy.blue/>.
- [16] URL: https://focus.jandenul.com/focus-magazine_11/seavolt.
- [17] D. Mehta, A.H. Van Zuijlen, B. Koren, J.G. Holierhoek, and H. Bijl. «Large Eddy Simulation of wind farm aerodynamics: A review». In: *Journal of Wind Engineering and Industrial Aerodynamics* 133 (Oct. 2014), pp. 1–17. ISSN: 01676105. DOI: 10.1016/j.jweia.2014.07.002.
- [18] S.-P. Breton, J. Sumner, J. N. Sørensen, K. S. Hansen, S. Sarmast, and S. Ivanell. «A survey of modelling methods for high-fidelity wind farm simulations using large eddy simulation». In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (Apr. 13, 2017), p. 20160097. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2016.0097.
- [19] URL: <https://www.exascaleproject.org/research-project/exawind/>.
- [20] URL: <https://github.com/Exawind/amr-wind>.
- [21] URL: <https://github.com/Exawind/nalu-wind>.
- [22] URL: <https://github.com/OpenFAST>.
- [23] W.J.M. Rankine. «On the Mechanical Principles of the Action of Propellers». In: (1865). *Transaction of the Institute of Naval Architects*, 6, 13-39.
- [24] Jens Nørkær Sørensen and Wen Zhong Shen. «Numerical Modeling of Wind Turbine Wakes». In: *Journal of Fluids Engineering* 124.2 (May 28, 2002), pp. 393–399. ISSN: 0098-2202. DOI: 10.1115/1.1471361.

-
- [25] David Maniaci and Alan Hsieh. *Wind Energy Physics-Based Modeling Through High-Performance Computing*. Tech. rep. Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2021.
- [26] Luis A Martinez Tossas and Stefano Leonardi. *Wind turbine modeling for computational fluid dynamics: December 2010-December 2012*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2013.
- [27] Misun Min, Michael Brazell, Ananias Tomboulides, Matthew Churchfield, Paul Fischer, and Michael Sprague. *Towards Exascale for Wind Energy Simulations*. Sept. 30, 2022. arXiv: 2210.00904[cs]. URL: <http://arxiv.org/abs/2210.00904>.
- [28] Wikipedia contributors. *Summit (supercomputer) — Wikipedia, The Free Encyclopedia*. [https://en.wikipedia.org/w/index.php?title=Summit_\(supercomputer\)&oldid=1241305710](https://en.wikipedia.org/w/index.php?title=Summit_(supercomputer)&oldid=1241305710). 2024.
- [29] Wikipedia contributors. *Frontier (supercomputer) — Wikipedia, The Free Encyclopedia*. [https://en.wikipedia.org/w/index.php?title=Frontier_\(supercomputer\)&oldid=1243927543](https://en.wikipedia.org/w/index.php?title=Frontier_(supercomputer)&oldid=1243927543). 2024.
- [30] Tanmoy Chatterjee, Jing Li, Shashank Yellapantula, Balaji Jayaraman, and Eliot Quon. «Wind farm response to mesoscale-driven coastal low level jets: a multiscale large eddy simulation study». In: *Journal of Physics: Conference Series*. Vol. 2265. 2. IOP Publishing. 2022, p. 022004.
- [31] Lawrence Cheung et al. «Investigations of Farm-to-Farm Interactions and Blockage Effects from AWAKEN Using Large-Scale Numerical Simulations». In: *Journal of Physics: Conference Series* 2505.1 (May 1, 2023), p. 012023. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/2505/1/012023.
- [32] URL: <https://www.nrel.gov/wind/awaken.html>.
- [33] Weiqun Zhang, Andrew Myers, Kevin Gott, Ann Almgren, and John Bell. «AMReX: Block-structured adaptive mesh refinement for multiphysics applications». In: *The International Journal of High Performance Computing Applications* 35.6 (Nov. 1, 2021). Publisher: SAGE Publications Ltd STM, pp. 508–526. ISSN: 1094-3420. DOI: 10.1177/10943420211022811.
- [34] Ashesh Sharma et al. «ExaWind: Open-source CFD for hybrid-RANS/LES geometry-resolved wind turbine simulations in atmospheric flows». In: *Wind Energy* 27.3 (2024), pp. 225–257.
- [35] J_ Boussinesq. «Theorie analytique de la chaleur vol 2 (paris: Gauthier-villars)». In: *Buoyancy Effects in Fluids* (1903).

- [36] Chin-Hoh Moeng. «A Large-Eddy-Simulation Model for the Study of Planetary Boundary-Layer Turbulence». In: *Journal of the Atmospheric Sciences* (July 1, 1984). ISSN: 1520-0469. DOI: 10.1175/1520-0469(1984)041%3C2052:ALESMF%3E2.0.CO;2.
- [37] Jean Boussinesq. *Essai sur la théorie des eaux courantes*. Imprimerie nationale, 1877.
- [38] Andrei Sergeevich Monin and Aleksandr Mikhailovich Obukhov. «Basic laws of turbulent mixing in the surface layer of the atmosphere». In: *Contrib. Geophys. Inst. Acad. Sci. USSR* 151.163 (1954), e187.
- [39] Ann S. Almgren, John B. Bell, Phillip Colella, Louis H. Howell, and Michael L. Welcome. «A Conservative Adaptive Projection Method for the Variable Density Incompressible Navier–Stokes Equations». In: *Journal of Computational Physics* 142.1 (May 1, 1998), pp. 1–46. ISSN: 0021-9991. DOI: 10.1006/jcph.1998.5890.
- [40] URL: <https://openfast.readthedocs.io/en/dev/source/user/aerodynamolaf/Introduction.html>.
- [41] Gene M Amdahl. «Validity of the single processor approach to achieving large scale computing capabilities». In: *Proceedings of the April 18-20, 1967, spring joint computer conference*. 1967, pp. 483–485.
- [42] John L Gustafson. «Reevaluating Amdahl’s law». In: *Communications of the ACM* 31.5 (1988), pp. 532–533.
- [43] URL: <https://energy.sandia.gov/programs/renewable-energy/wind-power/swift-facilities/>.
- [44] URL: <https://iea-wind.org/task31/>.
- [45] URL: <https://wakebench-swift.readthedocs.io/en/latest/index.html#>.
- [46] URL: <https://doi.org/10.1016/j.rser.2015.01.052>.
- [47] Alan S. Hsieh et al. «High-fidelity wind farm simulation methodology with experimental validation». In: *Journal of Wind Engineering and Industrial Aerodynamics* 218 (Nov. 2021), p. 104754. ISSN: 01676105. DOI: 10.1016/j.jweia.2021.104754.
- [48] TG Herges, David Charles Maniaci, Brian Thomas Naughton, Torben Mikkelsen and Mikael Sjöholm. «High resolution wind turbine wake measurements with a scanning lidar». In: *Journal of Physics: Conference Series*. Vol. 854. 1. IOP Publishing. 2017, p. 012021.

- [49] TG Herges and Patrick Keyantuo. «Robust lidar data processing and quality control methods developed for the SWiFT wake steering experiment». In: *Journal of Physics: Conference Series*. Vol. 1256. 1. IOP Publishing. 2019, p. 012005.
- [50] Alan Hsieh, David C. Maniaci, Thomas G. Herges, Gianluca Geraci, Daniel T. Seidl, Michael S. Eldred, Myra L. Blaylock, and Brent C. Houchens. «Multilevel Uncertainty Quantification Using CFD and OpenFAST Simulations of the SWiFT Facility». In: *AIAA Scitech 2020 Forum*. AIAA Scitech 2020 Forum. Orlando, FL: American Institute of Aeronautics and Astronautics, Jan. 6, 2020. ISBN: 978-1-62410-595-1. DOI: 10.2514/6.2020-1949.
- [51] Paula Doubrawa et al. «Multimodel validation of single wakes in neutral and stratified atmospheric conditions». In: *Wind Energy* 23.11 (Nov. 2020), pp. 2027–2055. ISSN: 1095-4244, 1099-1824. DOI: 10.1002/we.2543.
- [52] Niels Troldborg. «Actuator line modeling of wind turbine wakes». In: (2009).
- [53] URL: https://github.com/NREL/wakebench_swift.
- [54] Christopher Lee Kelley and Jonathan White. *An update to the SWiFT V27 reference model*. Tech. rep. Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2018.
- [55] Luis A Martínez-Tossas and Charles Meneveau. «Filtered lifting line theory and application to the actuator line model». In: *Journal of Fluid Mechanics* 863 (2019), pp. 269–292.
- [56] Myra L Blaylock, Luis Martinez-Tossas, Philip Sakievich, Brent C Houchens, Lawrence Cheung, Kenneth Brown, Alan Hsieh, David C Maniaci, and Matthew J Churchfield. «Validation of actuator line and actuator disk models with filtered lifting line corrections implemented in Nalu-Wind large eddy simulations of the atmospheric boundary layer». In: *AIAA SCITECH 2022 Forum*. 2022, p. 1921.
- [57] Fernando Porté-Agel, Majid Bastankhah, and Sina Shamsoddin. «Wind-turbine and wind-farm flows: a review». In: *Boundary-layer meteorology* 174.1 (2020), pp. 1–59.
- [58] Sina Shamsoddin and Fernando Porté-Agel. «Large-Eddy Simulation of Atmospheric Boundary-Layer Flow Through a Wind Farm Sited on Topography». In: *Boundary-Layer Meteorology* 163.1 (Apr. 2017), pp. 1–17. ISSN: 0006-8314, 1573-1472. DOI: 10.1007/s10546-016-0216-z.
- [59] Luoqin Liu and Richard J. A. M. Stevens. «Effects of Two-Dimensional Steep Hills on the Performance of Wind Turbines and Wind Farms». In: *Boundary-Layer Meteorology* 176.2 (Aug. 2020), pp. 251–269. ISSN: 0006-8314, 1573-1472. DOI: 10.1007/s10546-020-00522-z.

- [60] Charles S Peskin. «Flow patterns around heart valves: a numerical method». In: *Journal of computational physics* 10.2 (1972), pp. 252–271.
- [61] Stuart Chester, Charles Meneveau, and Marc B Parlange. «Modeling turbulent flow over fractal trees with renormalized numerical simulation». In: *Journal of Computational Physics* 225.1 (2007), pp. 427–448.
- [62] Mingwei Ge, Dennice F. Gayme, and Charles Meneveau. «Large-eddy simulation of wind turbines immersed in the wake of a cube-shaped building». In: *Renewable Energy* 163 (Jan. 2021), pp. 1063–1077. ISSN: 09601481. DOI: 10.1016/j.renene.2020.08.156.
- [63] Naveen N Kethavath and Niranjan S Ghaisas. «1 Effect of an abrupt rough-to-smooth surface roughness transition on wind 2 farm wakes: An LES and analytical modelling study». In: (Feb. 7, 2024).
- [64] Nicolas Tobin and Leonardo Chamorro. «Windbreak Effects Within Infinite Wind Farms». In: *Energies* 10.8 (Aug. 3, 2017), p. 1140. ISSN: 1996-1073. DOI: 10.3390/en10081140.
- [65] Luoqin Liu and Richard J. A. M. Stevens. «Enhanced wind-farm performance using windbreaks». In: *Physical Review Fluids* 6.7 (July 30, 2021), p. 074611. ISSN: 2469-990X. DOI: 10.1103/PhysRevFluids.6.074611.
- [66] Evan Gaertner et al. *IEA wind TCP task 37: definition of the IEA 15-megawatt offshore reference wind turbine*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020.
- [67] Marc Calaf, Charles Meneveau, and Johan Meyers. «Large eddy simulation study of fully developed wind-turbine array boundary layers». In: *Physics of Fluids* 22.1 (Jan. 1, 2010), p. 015110. ISSN: 1070-6631, 1089-7666. DOI: 10.1063/1.3291077.
- [68] Richard J.A.M. Stevens, Jason Graham, and Charles Meneveau. «A concurrent precursor inflow method for Large Eddy Simulations and applications to finite length wind farms». In: *Renewable Energy* 68 (Aug. 2014), pp. 46–50. ISSN: 09601481. DOI: 10.1016/j.renene.2014.01.024.
- [69] Tony Burton, Nick Jenkins, David Sharpe, and Ervin Bossanyi. *Wind energy handbook*. John Wiley & Sons, 2011.
- [70] URL: <https://github.com/IEAWindTask37/IEA-15-240-RWT/tree/master/OpenFAST/IEA-15-240-RWT-Monopile>.
- [71] Richard JAM Stevens, Dennice F Gayme, and Charles Meneveau. «Effects of turbine spacing on the power output of extended wind-farms». In: *Wind Energy* 19.2 (2016), pp. 359–370.