

Wake Modelling of a Floating Offshore Wind Turbine

Tuning of a Large Eddy Simulation with an Actuator Line Method for a FOWT with Sinusoidal Surge and Pitch Motion

Roel van Leeuwen

Wake Modelling of a Floating Offshore Wind Turbine

Tuning of a Large Eddy Simulation with an Actuator Line Method for a FOWT with Sinusoidal Surge and Pitch Motion

by

Roel van Leeuwen

to obtain the degree of Master of Science
at the Delft University of Technology.

Student number:	5257069	
Project duration:	December 9, 2024 - August 28, 2025	
Faculty:	Faculty of Aerospace Engineering, Delft	
Thesis committee:	Prof.dr.ir. A.C. Viré	Chair
	Dr. ir. A.H. van Zuijlen	Examiner
	Dr. F. Taruffi	Supervisor
	E. Wiegant	Co-supervisor

Preface

This thesis represents the final step of my studies and focuses on simulating the wake of a floating wind turbine. Conducting research in this area of wind energy has been a dream come true, as my fascination with the topic began three years ago during my minor in offshore wind energy.

Over the past year, I have grown to appreciate even more what a privilege it is to study at TU Delft. The abundance of resources and the stimulating environment have continuously motivated me to learn more and work harder.

Throughout the years, I have received support from many people. I am deeply grateful to my family and friends for their unconditional encouragement and for challenging me to keep improving in all aspects. I would also like to thank the team at MuTech B.V., where I completed my internship; their dedication and drive inspired me to push forward.

During my thesis work, I was fortunate to share an office with incredible colleagues, who made this period not only productive but also enjoyable. Their constructive feedback—both on the thesis and on personal matters—helped me grow, and many shared cookies later, I believe they have made me a better person overall.

Lastly, I would like to thank my supervisors, Federico and Evert, for their guidance, knowledge, and support throughout this journey.

*Roel van Leeuwen
Delft, August 2025*

Summary

In recent years, floating offshore wind turbines have emerged as a key technology for deploying wind energy in locations with water depths exceeding 50 meters, where bottom-fixed solutions become economically unfeasible [1]. These deeper waters offer a global potential of approximately 13 TW [1]. However, unlike fixed-bottom turbines, floating turbines exhibit lower structural stiffness due to their reliance on flexible mooring systems. This increased flexibility causes platform motion in response to wind and wave loads, introducing significant challenges in turbine and wind farm design. A particularly important challenge lies in understanding how rotor wake aerodynamics—critical for optimising farm layout and energy production—are influenced by such motion. While previous studies have shown that floating motion affects wake recovery, there is no consensus on whether it accelerates or slows this process.

The objective of this work is to numerically replicate the wake aerodynamics of a scale wind turbine model undergoing prescribed sinusoidal surge and pitch motion, using a Large-Eddy Simulation Actuator Line Method. At TU Delft, wind tunnel experiments were conducted on a scaled DTU 10 MW turbine model, where platform motion was imposed via a robotic motion platform. These experimental results are used to calibrate and validate the numerical framework for floating wind turbine wake modelling, providing a rare opportunity for optimisation given the scarcity of such experimental datasets. A validated numerical model can complement the experiments by identifying relevant motion cases for further testing and by correcting experimental results for effects such as wind tunnel interference.

The simulation framework, known as AspFAST, couples the Large-Eddy Simulation code GPU-Resident Atmospheric Simulation Platform (GRASP) with the multi-physics, multi-fidelity tool OpenFAST. It offers a broad range of configurable numerical settings. A preliminary sensitivity analysis was carried out to examine the influence of key modelling parameters. Based on these results, a baseline setup was established for both a stationary case and four dynamic motion cases: two surge motions and two pitch motions, each with different frequencies.

The simulations were assessed by comparing time-averaged flow quantities—specifically, the streamwise velocity, out-of-plane vorticity, turbulence intensity, and streamwise–vertical Reynolds stress—with corresponding wind tunnel measurements. To more accurately replicate the experimental environment, a turbulent inflow was subsequently introduced, which enhanced wake instability and led to better alignment with the experimental results. Nevertheless, the simulated wake effects induced by platform motion remained considerably less distinct than those observed in the experiments.

Further investigations examined the impact of several model refinements: the Filtered Actuator Line Method, the inclusion of nacelle, and increased mesh resolution. While the Filtered Actuator Line Method did not yield noticeable improvements, including the nacelle led to marginally better modelling of the near wake in the root region. Increasing the resolution improved representation of the tip vortex significantly.

Contents

Preface	i
Summary	ii
List of figures	x
List of tables	xi
List of symbols	xi
I Introduction & Theory	1
1 Introduction	2
1.1 Floating Wind Turbines	2
1.2 Scope and Research Questions	3
1.3 Structure	4
2 Theory & Literature	5
2.1 Rotor Aerodynamics	5
2.2 Wake Aerodynamics	8
2.3 Numerical Methods for Testing FOWTs	9
2.4 Experimental Methods for Testing FOWTs	11
2.5 Numerical–Experimental Evaluation of FOWT Aerodynamics	13
II Methodology	16
3 Experimental Setup	17
3.1 Test Environment	17
3.2 Wind Turbine Scale Model	17
3.3 Motion Cases	18
3.4 Particle Tracking Velocimetry	18
4 Numerical Setup	21
4.1 GRASP: Large Eddy Simulation	21
4.2 OpenFAST: Wind Turbine Characterisation	26
4.3 AspFAST: Actuator Line Method	28
4.4 Hardware	31
4.5 Numerical-Experimental Comparison	31
III Simulation Tuning	33
5 Parameter Study	34
5.1 Inflow and Outflow Boundary condition	34
5.2 Domain Size	34
5.3 Actuator Points	37
5.4 Kernel Size	37
5.5 Grid Resolution	38
5.6 Base Case	40

5.7	FALM	40
5.8	Turbulence Synthesis	41
6	Experimental Results	44
6.1	Streamwise Velocity	44
6.2	Vorticity	45
6.3	Turbulence Intensity	47
6.4	Reynolds Stress	49
6.5	Vertical Velocity	50
6.6	Linking the Metrics	50
7	Base Simulation	54
7.1	Streamwise Velocity	54
7.2	Vorticity	55
7.3	Turbulence Intensity	57
7.4	Reynolds stress	58
7.5	Linking the Metrics	60
8	Turbulent Inflow	62
8.1	Turbulence Generation	62
8.2	Streamwise Velocity	63
8.3	Vorticity	65
8.4	Turbulence intensity	66
8.5	Reynolds stress	68
8.6	Vertical Velocity	70
8.7	Linking the Metrics	71
9	FALM, Nacelle and Resolution	74
9.1	FALM	74
9.2	Including Nacelle	74
10	High Resolution	77
10.1	Inflow	77
10.2	Vertical Velocity	77
10.3	Streamwise Velocity	77
10.4	Vorticity	79
10.5	Turbulence Intensity	80
10.6	Linking the Metrics	81
IV	Conclusion	83
11	Conclusions and Future Scope	84
11.1	Conclusions	84
11.2	Future Scope	85
	Appendix	93
A	Simulation tuning	94
A.1	Experimental results	94
A.2	Base simulation	99
A.3	Turbulent inflow	103
A.4	FALM	108
A.5	Including Nacelle	109

A.6	High resolution	110
B	Sub-Grid Scale Model and Advection Scheme Choice	114
B.1	Turbulence Spectrum	114
B.2	Streamwise Velocity	115
B.3	Reynolds stress	115
C	Tower Modelling	118
C.1	Streamwise Velocity	118
C.2	Vorticity	118
C.3	Turbulence Intensity	119
C.4	Reynolds stress	119
D	Vertical Shift Correction	122
D.1	Correction Method	122
D.2	Vertical Shift Results	122

List of Figures

1.1	Coordinate system and Degrees of Freedom (DOFs) of a floating wind turbine [2]	3
2.1	Schematic of a blade section, including geometric and force terms [3].	5
2.2	Schematic of the vortex filament sheet behind a blade [4]	7
2.3	Schematic of wake evolution from a top-hat profile to a Gaussian profile. The mixing layer growth is visible as wake expansion and decrease of potential core size [5].	8
3.1	Schematic of the Particle Tracking Velocimetry (PTV) setup [6].	19
4.1	Schematic of the numerical framework (modified from [7]). The central coupling module AspFAST integrates the Large Eddy Simulation (LES) solver GRASP, developed by Whiffle, via the Atmospheric Simulation Platform for Innovation Research and Education (ASPIRE) interface (left), with OpenFAST (right). For this study, the Aerodynamics, Elastodynamics, and Sub-structure dynamics modules of OpenFAST are used.	21
4.2	Schematic of the pitch motion with amplitude A_p . Due to the extended tower height from H_{t_0} to H_t , the motion must be compensated with a surge and heave oscillation with amplitudes A_s and A_h	28
4.3	Contours of the summed force distribution kernel η for $\epsilon/\Delta x$ equal to 2 (left) and 5 (right). The blue dots are the actuator points of the actuator lines.	29
5.1	Normalised, time-averaged streamwise velocity at the end of the upwind nudging zone. The streamwise nudging extent is varied from 6 to 20 cells. The nudging factor is equal to 1 in all simulations.	35
5.2	Normalised, time-averaged, rotor-averaged streamwise velocity, upwind of the rotor. The five curves exclude the upwind nudging zone and show that the upwind velocity converge to one velocity development as the upwind distance increases.	36
5.3	Normalised, time-averaged, rotor-averaged streamwise velocity, downwind of the rotor. The four curves exclude the downwind nudging zone and show that approximately 2 D of downwind domain is unphysical due to the influence of the outflow boundary.	36
5.4	Thrust coefficient C_T against the blockage ratio, indicating a linear reduction of thrust as blockage reduces.	37
5.5	Normalised, time-averaged velocity along the vertical direction at $x/D = 2$. The velocity converges for low blockage ratios.	37
5.6	Normalised, time-average streamwise velocity profile along the vertical direction at $x/D = 2$ for different numbers of actuator points. The profile converges for 50 actuator points on each blade.	38
5.7	Effect of reducing the kernel size ϵ on the normalised vertical velocity profile at $x/D = 2$	38
5.8	Fraction of resolved Turbulent Kinetic Energy (TKE) with respect to the total TKE in the near wake for varying mesh resolution. The resolutions $\Delta x/D$ from top to bottom are 0.017, 0.012 and 0.011	39
5.9	Time-averaged, space-averaged ratio of resolved and total TKE averaged from the turbine rotor plane to the domain outflow and between $z/D = \pm 1$ along the vertical direction. The trend through the measurements shows an exponential fit.	40
5.10	Normalised, time-averaged streamwise velocity along the vertical direction at $x/D = 2$ for varying resolutions. A higher resolution increases the velocity gradients, but could cause an insufficient kernel size.	40

5.11	Schematic of the domain. The white area is physically useful, while the gray area is affected by the boundary conditions. Left is a front view and right is a side view. The turbine point of rotation is indicated by the black bullet. The red arrow indicates the surge motion and the blue arrow indicates the pitch motion.	41
5.12	Kernel size with respect to the local chord length for two kernel sizes. Both have an $\epsilon/\Delta x$ of 2, but with resolutions $\Delta x/D$ of 0.041 and 0.012.	42
5.13	Effect of the Filtered Actuator Line Method (FALM) on the normalised, time-averaged streamwise velocity. The blue curves have a relative kernel size of 2 and the green curves one of 5, with equal resolutions.	42
5.14	Effect of heat perturbation magnitude and heat flux on the Turbulence Intensity (TI). . .	43
6.1	Normalised time-averaged streamwise velocity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	46
6.2	Normalised time-averaged out-of-plane vorticity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	47
6.3	Time-averaged TI along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	49
6.4	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	51
6.5	Normalised time-averaged vertical velocity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	52
7.1	Normalised time-averaged streamwise velocity along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	56
7.2	Normalised time-averaged out-of-plane vorticity along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	57
7.3	Time-averaged TI along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	59
7.4	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	60
8.1	Turbulent inflow characteristics.	62
8.2	Resolved TKE as a fraction of total TKE for the stationary turbine with a uniform 3.6% TI inflow.	63
8.3	Normalised time-averaged streamwise velocity along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	65
8.4	Normalised time-averaged out-of-plane vorticity along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	67

8.5	Time-averaged TI along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	69
8.6	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	71
8.7	Normalised time-averaged vertical velocity along the vertical direction from the base simulation (solid line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	72
9.1	Normalised time-averaged out-of-plane vorticity difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the Actuator Line Method (ALM) and experiments are presented (dotted line). All differences are along the vertical direction of the domain for all motion cases. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.	74
9.2	Normalised time-averaged streamwise velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.	75
9.3	Normalised time-averaged out-of-plane velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.	76
10.1	High-resolution precursor inflow conditions compared to the inflow conditions from the low-resolution precursor. Top-left is the time-averaged TI (solid lines) and its standard deviation (dotted lines). Top-right is the normalised time-averaged vertical velocity. Bottom are the spatial turbulence spectra.	78
10.2	Normalised time-averaged vertical velocity along the vertical direction from the base simulation (solid line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).	79
10.3	Normalised time-averaged streamwise velocity along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	80
10.4	Normalised time-averaged out-of-plane vorticity along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.	81
10.5	Time-averaged TI along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.	82
10.6	Time-averaged thrust per unit volume along the actuator line for the high-resolution simulations (solid lines) and low-resolution simulations (dashed lines).	82
A.1	Normalised time-averaged streamwise velocity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	94
A.2	Normalised time-averaged out-of-plane vorticity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	95
A.3	Time-averaged TI field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	96

A.4	Normalised time-averaged streamwise-vertical Reynolds stress field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	97
A.5	Normalised time-averaged vertical velocity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	98
A.6	Normalised time-averaged streamwise velocity field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	99
A.7	Normalised time-averaged streamwise velocity field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	100
A.8	Time-averaged TI field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	101
A.9	Normalised time-averaged streamwise-vertical Reynolds stress field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	102
A.10	Normalised time-averaged streamwise velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	103
A.11	Normalised time-averaged streamwise velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	104
A.12	Time-averaged TI field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	105
A.13	Normalised time-averaged streamwise-vertical Reynolds stress field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	106
A.14	Normalised time-averaged vertical velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.	107
A.15	Normalised time-averaged streamwise velocity difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.	108
A.16	Time-averaged TI difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.	108
A.17	Normalised time-averaged streamwise-vertical Reynolds stress difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.	109
A.18	Time-averaged TI along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to . 0.7, 1.3, 1.9 and 2.5.	109

A.19	Normalised time-averaged out-of-plane velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.	110
A.20	Normalised time-averaged streamwise velocity along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	110
A.21	Normalised time-averaged out-of-plane vorticity along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	111
A.22	Time-averaged TI along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	112
A.23	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.	113
B.1	Spatial turbulence spectra for the turbulent inflow simulation setup (see Chapter 8) with the Smagorinsky Subgrid Scale (SGS) model and the Verstappen model with a second and fifth order advection scheme. The grey line indicates the -5/3 slope in the inertial subrange.	115
B.2	Normalised time-averaged streamwise velocity along the vertical direction from the simulation with turbulent inflow. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. In addition, the experimental results for the stationary case have been added as the grey dashed line.	116
B.3	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the simulation with turbulent inflow. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. In addition, the experimental results for the stationary case have been added as the grey dashed line.	117
C.1	Normalised time-averaged streamwise velocity along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.	119
C.2	Normalised time-averaged out-of-plane vorticity along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.	120
C.3	Time-averaged TI along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.	120
C.4	Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.	121
D.1	Block diagram illustrating the method for applying a vertical position correction to the simulated wake relative to the experiments.	122
D.2	Normalised time-averaged streamwise velocity along the vertical direction for simulations including a nacelle, after applying the vertical shift correction. Profiles are shown at $x/D = 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4$ and 5.1.	123

List of Tables

3.1	Wind turbine scale model design parameters.	18
3.2	Motion cases test matrix.	18

List of Symbols

Upper-Case Roman

A	Area	F_N	Normal force
A_h	Heave amplitude	F_{sub}	Substructure loads
A_p	Pitch amplitude	F_T	Tangential force
A_s	Surge amplitude	Fr	Froude number
A_θ	Heat perturbation magnitude	K	Stiffness
C	Damping	L	Length
C_d	Airfoil drag coefficient	M	Mass
C_L	Blade lift coefficient	N_b	Number of actuator points
C_l	Airfoil lift coefficient	N_x	Number of nudging cells
C_P	Power coefficient	P	Power
C_T	Thrust coefficient	Q	Torque
C_x	Subgrid scale model constant	R	Radius
D	Rotor diameter	Re	Reynolds number
D_t	Tower diameter	S	Strain rate tensor
F_{aero}	Aerodynamic loads	U	Fluid velocity
F_D	Drag force	U_0	Reference flow velocity
F_L	Lift force	U_∞	Free stream velocity
		U_{rel}	Relative fluid velocity
		V	Platform velocity

Lower-Case Roman

a	Induction	r	Radial position
c	Chord length	s	Scalar
c_s	Constant	t	Time
f	Frequency	u	Streamwise velocity (in the x-direction)
f_0	Natural frequency	u_y	Downwash velocity
f_e	Distributed force	v	Lateral velocity (in the y-direction)
f_{AL}	Actuator Line force	w	Vertical velocity (in the z-direction)
g	Gravitational acceleration	w_f	Force weighting
k	Turbulent kinetic energy	x, y, z	Coordinates
p	Pressure	z_0	Roughness length
q	Floater state		

Upper-Case Greek

Δ	Filter width
ΔV^*	Normalised hub velocity amplitude
Δx	Cell size
Γ	Circulation
Ω	Azimuthal velocity

Lower-Case Greek

α	Angle of attack	ν	Specific kinematic viscosity
α_n	Relaxation factor	ϕ	Twist
δx	Spatial step	ρ	Density
ϵ	Dissipation / kernel size	θ	Rotation
η	Distribution weight	$\theta_{p,0}$	Global pitch angle
η_n	Nudging factor	θ_T	Potential temperature
λ	Scaling factor	$\bar{\omega}$	Model operator

Subscripts

rms	Root mean square
x, y, z	Space dimensions

Symbols

$\bar{\cdot}$	Spatial average
$\langle \cdot \rangle$	Time average
$\tilde{\cdot}$	Filtered variable

Part I

Introduction & Theory

Introduction

Wind turbines have become one of the key technologies enabling the decarbonisation of the global energy system. Although wind energy has been harnessed since prehistoric times [3], the modern development of wind turbines gained momentum in the early 20th century, driven by the increasing demand for electricity. The urgency of reducing greenhouse gas emissions further accelerated their advancement in the late 20th century, as the environmental consequences of fossil fuel combustion became widely recognised.

Initial wind turbine deployments were land-based, but the greater wind resource available offshore soon attracted attention. This led to the commissioning of the world's first offshore wind farm, Vindeby, off the coast of Denmark in 1991. Since then, numerous offshore wind farms have been developed, primarily using bottom-fixed turbines located in shallow coastal waters. Such installations are generally limited to depths of less than 50 metres due to economic constraints [1].

However, the wind resource increases with distance from the shore, and the availability of suitable shallow-water sites is diminishing—many have already been developed or are reserved for future use, whether for wind energy or other purposes. To access deeper waters, where an estimated 13 TW of technical potential exists [1], Floating Offshore Wind Turbines (FOWTs) offer a promising solution. The concept dates back to 1972, when [8] first proposed floating wind systems. In recent years, interest in FOWTs has grown significantly as a means of overcoming the depth limitations of bottom-fixed designs.

1.1. Floating Wind Turbines

FOWTs differ from fixed-bottom turbines primarily in their dynamic behaviour. Rather than being rigidly connected to the seabed, they are mounted on a floater that provides buoyancy and stability, and are held in position using mooring lines and anchors. This configuration results in a relatively flexible substructure compared to bottom-fixed turbines, while operating in a highly energetic and challenging offshore environment. Ocean waves, which can reach several metres in height and carry substantial energy, impose dynamic loads on the structure. Additionally, the wind exerts aerodynamic forces on the turbine. The combined effects of wind and wave loading, together with the mass, stiffness, and damping of the FOWT, cause it to oscillate in six DOFs.

These DOFs comprise translations along and rotations about the axes of an inertial coordinate system, as illustrated in Figure 1.1 [2]. The x -axis is aligned with the mean wind direction and is positive in the downwind direction. Translation along this axis is referred to as *surge*, while rotation about it is called *roll*. The positive y -axis points towards the starboard side of the turbine, with *sway* and *pitch* denoting translation and rotation, respectively. The z -axis is oriented vertically along the tower, pointing towards the nacelle. Motion along this axis is known as *heave*, and rotation about it as *yaw*.

The oscillatory motion in these DOFs leads to time-varying relative velocities at the rotor, which influence the aerodynamic behaviour of the FOWT. These velocities can be described by:

$$[V_x \ V_y \ V_z] = [\dot{x} \ \dot{y} \ \dot{z}] + [\dot{\theta}_x \ \dot{\theta}_y \ \dot{\theta}_z] \times [x \ y \ z] \quad (1.1)$$

Here, \dot{x} , \dot{y} , \dot{z} represent translational velocities, and $\dot{\theta}_x$, $\dot{\theta}_y$, $\dot{\theta}_z$ denote angular velocities about the respective axes. The cross-product term accounts for the contribution of rotational motion to the velocity field at a given point on the rotor.

[9] were the first to address this fundamental difference with the conventional bottom-fixed turbine. Since then, the aerodynamics of FOWTs has been researched both experimentally and numerically. It has been




Figure 1.1: Coordinate system and DOFs of a floating wind turbine [2]

found that in terms of turbine performance not all DOFs are equally relevant. It was shown that surge, pitch and/or roll contain most of the environment-induced platform motion [10]. Besides, the flow normal to the rotor is what transfers momentum to the rotor. Roll causes opposite motion induced velocity at the port and starboard side of the rotor, eliminating a significant change in performance. For these reasons most research has been conducted for turbines in surge and pitch motion.

1.2. Scope and Research Questions

At the Delft University of Technology, a robotic motion platform is used inside an open-jet wind tunnel to test a scale-model of the DTU 10MW wind turbine [11] with mimicked motion due to waves and wind loading, or with prescribed harmonic motion at distinct frequencies. This allows for performing experiments with highly accurate performance and flow results [12, 13, 14]. In literature, there is no consensus on whether the turbine motion accelerates or slows the wake recovery. Recently, wake measurements have been conducted for turbines undergoing surge and pitch motions at various frequencies and amplitudes [6], aiming to improve understanding of wake aerodynamics under different floater motions.

To complement these experiments, a validated numerical version of the experiments is desired to complement the experiments. For this purpose, Whiffle Weather Finecasting Ltd¹ has provided their in-house Large Eddy Simulation (LES), called GPU-Resident Atmospheric Simulation Platform (GRASP), which is coupled through AspFAST to the multi-physics, multi-fidelity tool, called OpenFAST. This study describes the process of identifying an appropriate simulation configuration in pursuit of the following aim:

To numerically replicate the **wake aerodynamics** of a scale wind turbine model with **prescribed sinusoidal surge** and **pitch** motion, using a **Large Eddy Simulation (LES) Actuator Line Method (ALM)** in AspFAST.

This aim is addressed through the following research questions, which break down the problem into manageable components:

¹<https://whiffle.nl/>

1. Which configuration options in AspFAST affect the wake development, and how?
2. What are the characteristics of the experimentally measured wake, and how do they differ between motion cases?
3. How does inflow turbulence affect the wake?
4. How does the Filtered Actuator Line Method (FALM) affect the wake?
5. How does the nacelle drag affect the wake?

1.3. Structure

This report is divided into four parts. Part I introduces the research topic and states the main aim. Additionally, Chapter 2 provides the theoretical background on wind turbine aerodynamics, the state-of-the-art numerical methods for simulating the flow around operating wind turbines, and current experimental techniques for testing FOWTs. It also discusses how the rotor aerodynamics and performance change when the rotor moves due to floater motion.

Part II details the methodology employed in this research. First, Chapter 3 describes the experimental setup from [6], as this forms the basis for the numerical simulations. Then, Chapter 4 explains the numerical tools used and the key parameters involved in configuring the simulations.

Part III begins with a parameter analysis in Chapter 5, examining simulation convergence and behaviour of the wake with respect to various parameters, addressing research question 1. This knowledge informs the definition of the base simulation, whose results are discussed in Chapter 7 and compared against experimental findings in Chapter 6, thereby addressing research question 2. Subsequently, an appropriate level of turbulence is introduced into the inflow in Chapter 8 to investigate its effect on the wake, responding to question 3. This is followed by evaluation of the FALM and the inclusion of nacelle drag effect on the wake, which answer questions 4 and 5. Next, the mesh resolution is increased in Chapter 10 to improve the tip vorticity modelling in the near wake to aid the general research aim of this work.

Part IV concludes the report with conclusions, and an outline of the future scope.

Theory & Literature

This chapter provides a review of key concepts in wind turbine aerodynamics and summarises relevant research in the field of FOWT aerodynamics. First, the rotor aerodynamics of a conventional fixed-bottom wind turbine is outlined in Section 2.1, to support understanding of how the wake is generated by the rotor. Next, the characteristics of the wake behind a fixed-bottom turbine are discussed in Section 2.2.

Following this, Section 2.3 presents four commonly used methods for simulating the wake in current research. An overview of experimental approaches to testing FOWTs is then given in Section 2.4. Finally, the influence of floater motion on rotor and wake aerodynamics is examined for surge and pitch motion in Section 2.5.

2.1. Rotor Aerodynamics

The wake of a wind turbine arises from the extraction of momentum from the flow by the blades, nacelle, and tower. Among these components, the blades are primarily responsible for power generation and have the dominant influence on the resulting wake dynamics.

2.1.1. Blade Element Forces

Figure 2.1 (adopted from [3]) presents a schematic of the flow at a rotor blade section, illustrating the key parameters that influence rotor aerodynamics. This sectional view forms the basis of blade element theory, in which the aerodynamic performance of a blade is calculated by dividing it into discrete elements along the span.




Figure 2.1: Schematic of a blade section, including geometric and force terms [3].

In the schematic, the rotor experiences an inflow velocity $U_\infty(1-a)$, where U_∞ is the freestream velocity and a is the axial induction factor, which accounts for the reduction in axial velocity due to the thrust exerted by the rotor.

$$a = \frac{U_\infty - U_{rot}}{U_\infty} \quad (2.1)$$

In addition to the axial inflow, the blade section also experiences a rotational velocity. This consists of the sum of the blade section velocity Ωr and the induced angular velocity $\omega r/2$. The induced angular velocity is a result of conservation of angular momentum [3]. ω is an equivalent angular velocity in the plane of the blade section. This angular velocity is located somewhere in the flow and has a magnitude such that it combines all angular momentum present around the rotor. The rotation arm r is the distance from the rotational centre to the blade section. The ratio between the induced and rotor angular velocity defines the azimuthal induction factor a'

The vector sum of the axial and azimuthal velocity at the blade section forms the relative velocity U_{rel} under an angle of attack α . Neglecting three-dimensional flow effects, the lift and drag on a blade section can be calculated using Equation 2.2 and Equation 2.3, respectively.

$$dF_L = \frac{1}{2} \rho U_{rel}^2 C_l c dr \quad (2.2)$$

$$dF_D = \frac{1}{2} \rho U_{rel}^2 C_d c dr \quad (2.3)$$

where ρ , C_l , C_d , c and dr are the air density, lift coefficient, drag coefficient, local chord length and section width. First, the angle of attack is used in the lift and drag polars of the airfoil at the section to get the lift and drag coefficients, resulting in the lift and drag of the section. Note that all parameter in these equations, except the air density, depend on the radial position r , resulting in a distributed varying load from root to tip.

Using the sectional twist angle ϕ , the lift and drag can be decomposed into components along the axial and azimuthal directions of the blade. The section normal force dF_N is defined in Equation 2.4, and the section tangential force dF_T is defined in Equation 2.5.

$$dF_N = dF_L \cos \phi + dF_D \sin \phi \quad (2.4)$$

$$dF_T = dF_L \sin \phi - dF_D \cos \phi \quad (2.5)$$

2.1.2. 3D Effects

Blade element theory describes the aerodynamic loads on a blade as a series of two-dimensional sections, which can be integrated to obtain the full spanwise load distribution. However, a wind turbine blade is a three-dimensional object and therefore experiences three-dimensional flow effects that alter this distribution. The most significant of these effects are tip loss and stall delay.

Tip Loss Tip losses arise due to the spanwise pressure gradient near the blade tip, which drives flow from the pressure side to the suction side. This results in the formation of a tip vortex that induces downwash, thereby reducing the effective angle of attack and lift near the tip [15].

Stall Delay Centrifugal forces acting on the boundary layer induce a radial velocity component, which delays the onset of flow separation along the blade [16]. As a result, stall is postponed, effectively shifting its onset to higher angles of attack. This effect is non-uniform along the blade span: the inboard sections experience a more pronounced delay, while the outboard sections are largely unaffected. Moreover, the effect intensifies at higher Tip Speed Ratio (TSR), making it particularly relevant at rated wind speed. The Viterna method [17] and Du–Selig method [16] were developed to correct two-dimensional airfoil polars for these rotational effects.

2.1.3. Circulation and Vorticity

The wake aerodynamics is linked to the rotor aerodynamics via the blade circulation Γ , which, according to the Kutta-Joukowski theorem, is directly related to the blade lift distribution, as shown in Equation 2.6 [15]. This lift distribution generates a corresponding circulation distribution along the blade,

referred as *bound vorticity*. The circulation induces angular velocity around the associated vortex filaments.

$$dF_L = \rho U_{rel} d\Gamma \quad (2.6)$$

Helmholtz's theorems state that a vortex filament cannot terminate within a fluid and that its strength remains constant along its length [15]. In the context of varying lift along the blade span, these constraints imply that the bound vorticity must either form closed loops or be connected to trailing filaments that extend downstream to infinity, each maintaining constant circulation.

Consequently, the bound vorticity is discretised into small spanwise segments—analogue to blade element theory—each with a defined circulation. At the ends of these segments, *trailing vortex* filaments are introduced. These are closed downstream by *shed vorticity*, which arises from temporal variations in the bound circulation. The result is a vortex sheet composed of trailing and shed vorticity extending behind the blade, as illustrated in Figure 2.2 [4].




Figure 2.2: Schematic of the vortex filament sheet behind a blade [4]

If the lift were constant along the blade span, the bound vorticity would also be constant, resulting in trailing vorticity only at the blade root and tip. This highlights that a low lift gradient produces minimal shed vorticity. However, due to tip loss effects, lift decreases rapidly near the blade tip. Similarly, near the root, the reduced rotational velocity leads to a decrease in lift. Consequently, strong trailing vorticity is concentrated at both the root and tip regions.

2.1.4. Unsteady Aerodynamics

When the angle of attack at the rotor varies rapidly, unsteady aerodynamic effects become significant. These effects can be categorised into unsteady profile aerodynamics and dynamic inflow [18]. In unsteady profile aerodynamics, flow inertia becomes relevant: a change in angle of attack alters the pressure gradient in the blade's boundary layer, causing local flow acceleration. As a result, aerodynamic loads lag behind the instantaneous angle of attack due to the finite response time of the flow.

A distinct unsteady profile effect is dynamic stall. When the angle of attack oscillates rapidly and exceeds the static stall angle of attack, a leading edge vortex can be formed. At some point, this vortex sheds from the leading edge and travels as a low pressure region over the top of the blade, increasing the load temporarily.

Dynamic inflow refers to the transient response of the induced velocity field at the rotor to changes in aerodynamic loading. Variations in shed vorticity alter the induced velocities experienced by the blades, causing a lag between the change in loading and the resulting wake development.

2.2. Wake Aerodynamics

The wake begins downstream of the rotor, nacelle, and tower, and can be divided into distinct regions, each characterised by different physical processes. This section first discusses the wake structure for a stationary turbine, followed by descriptions of the several states the wake can be in.

2.2.1. Wake Regions

The wake can be divided in three downstream stages which are described by [19] and [20], and presented in Figure 2.3. It begins with the expansion region, where the initially irregular flow transitions toward a statistically uniform time-averaged velocity field. Inside the near wake vortex instabilities appear causing the eventual breakdown of vortex structures, leading to the generation of the far wake.




Figure 2.3: Schematic of wake evolution from a top-hat profile to a Gaussian profile. The mixing layer growth is visible as wake expansion and decrease of potential core size [5].

Expansion Region

The first one is the expansion region, lasting approximately 1 rotor diameter leeward of the rotor. In this region, the mean total and static pressure have dropped due to the momentum extraction at the rotor. Since momentum is extracted at the blades and not by the open parts of the rotor disk, the pressure and velocity distributions are not uniform.

After passing the rotor, the wake mixes and evolves towards a more uniform pressure and velocity distribution. During this process, the wake velocity decreases further. Although total pressure does not recover, static pressure increases while dynamic pressure decreases, consistent with Bernoulli's principle. To conserve mass, the wake expands until the pressure inside and outside the wake are identical.

Concurrently, vorticity shed from the blades remains structured in the form of trailing and shed vorticity, as discussed in Section 2.1.3. The wake induction causes the trailing vorticity to roll up near the blade tips and roots, forming concentrated helical tip and root vortices. As a result, most vorticity resides near these regions, with minimal vorticity in between. The induced velocities from these vortex structures contribute to the observed velocity deficit in the wake centre.

Near Wake

In the near wake, tip and root vortices dominate the wake structure. The root vortices typically break down shortly after the rotor due to strong interaction with the nacelle [21], while the tip vortices remain coherent over a longer downstream distance, marking the boundary between the low-speed wake interior and the high-speed freestream. The tip vortices exhibit instabilities that grow downstream, eventually

leading to vortex breakdown and transition to the far wake [22, 23, 21, 24].

Wake mixing is governed by both periodic coherent structures and random turbulent fluctuations. In low-turbulence inflow, the coherent vortices initially inhibit entrainment of momentum from the surrounding flow [21]. Phase-averaged measurements reveal that tip vortices act as a shielding mechanism, limiting energy exchange across the wake boundary. Their contribution to wake mixing is largely restricted to slow viscous diffusion. Once these vortices destabilise, random fluctuations dominate, and entrainment and turbulence production increase sharply [24].

Several instability modes have been identified for helical vortices, including short-wave, long-wave, and mutual-inductance instabilities [22]. The latter occurs due to variations in vortex spacing caused by local pitch variations, which enhances induced velocities between segments of the vortex system. [23] analysed the effect of the phase of neighbouring vortex filaments. It was found that the instabilities grow fastest when the neighbouring filaments are out-of-phase. Initially, the vortex instabilities grow linearly, independently of the harmonic perturbation amplitude. However, as the mutual-inductance increases, the instability becomes non-linear. At this point a typical instability occurs, called leapfrogging/vortex-pairing instability where two vortices rotate around each other and heavily interact before breaking down. Additionally, it was shown that vortex breakdown will occur earlier in the wake for higher inflow TI. [24] showed that the moment when the vortex pair has rotated 90° relative to each other, the wake mixing due to random fluctuations increases significantly.

Far Wake

The far wake begins once the vortex structures have fully broken down. For a stationary turbine, the resulting flow resembles an axisymmetric turbulent wake [25], where statistical variations occur only in the streamwise and radial directions.

The time-averaged velocity adopts a Gaussian profile, and turbulence production is minimal due to the weak velocity gradients. Instead, the dominant mechanism in the turbulent kinetic energy (TKE) budget is convection. Consequently, the far wake retains memory of its upstream history, particularly the conditions in the near wake.

2.2.2. Operating States

The operating state of a wind turbine wake is governed by the axial induction factor at the rotor, a , as outlined by [26]. Under normal operating conditions ($a < 0.5$), the turbine functions in the windmill brake state, where one-dimensional momentum theory remains valid. For negative values of a , the turbine enters the propeller state, acting as a thrust-producing rotor.

As a exceeds 0.5, the rotor enters the turbulent wake state. When $a \approx 1.0$, a local flow reversal develops near the hub and progresses radially outward. At $a = 1$, the axial flow through the rotor is entirely blocked, though flow continues outside the rotor streamtube. This flow regime is associated with a toroidal vortex structure, analogous to the Vortex Ring State (VRS) observed in helicopter aerodynamics [4]. For $a > 1$, the rotor operates in the propeller brake state, injecting energy into the flow.

[27] points out that the wake state is a property of the streamtube and not of the rotor. When the turbine experiences a change in thrust time is required for the force to accelerate the flow (negative or positive). Therefore, in a non-inertial reference system, i.e. no acceleration, a certain thrust coefficient does not guarantee a certain wake state.

2.3. Numerical Methods for Testing FOWTs

In the wind turbine aerodynamics field four numerical methods are commonly used for simulating the flow field around a (floating) wind turbine. The most common method is Blade Element Momentum Theory (BEMT) method which is a simple low-fidelity method, discussed in Section 2.3.1. The Vortex Wake methods in Section 2.3.2 provides a higher fidelity by simulating the vorticity on the blades and in the wake using Lifting Line Theory. The remaining methods are based on Computational Fluid Dynamics (CFD). These include actuator-based models in Section 2.3.3 and blade-resolved simulations in

Section 2.3.4. The highest fidelity is obtained with blade-resolved CFD where the blade boundary layer is resolved. A detailed review on these methods has been provided in [28].

2.3.1. Blade Element Momentum Theory Method

The BEMT method [29] combines Blade Element theory with One-dimensional Momentum theory. The former one is discussed previously in Section 2.1. It provides a way to calculate the distributed aerodynamic load by means of the blade geometry, freestream flow and the aerodynamic lift and drag polar of the airfoil profiles contained by the blade.

One-dimensional Momentum theory takes the stream tube through the rotor as control volume. Only flow enters through the upstream far-field inlet and leaves through the downwind far-field outlet. The rotor is modelled as an actuator disk applying a uniform thrust on the flow. The one-dimensional mass and momentum conservation is applied on the flow inside the stream tube.

In BEMT, this result is applied on independent stream tube annuli. Each annuli can have a different thrust like a real blade would have, leading to a local induction for that annulus. The induction is used in the Blade Element theory to obtain the inflow angle, angle of attack and aerodynamic coefficients at the annulus. This leads to a new thrust coefficient which relates to another induction. Through, under-relaxation the induction is correct and these steps are repeated until the induction from the Blade Element theory and 1D Momentum theory are the same.

This method uses two-dimensional polars and assumes an infinite number of blades. Therefore, it inherently neglects three-dimensional flow phenomena such as stall delay and tip vortex formation. Various corrections exist such as the tip correction to model a finite number of blades.

2.3.2. Vortex Wake Methods

The Vortex Wake Methods use the Lifting Line theory to model bound vorticity on the blade, and shed and trailing vorticity in the wake as vortex filaments. The early developed methods were for helicopter rotor aerodynamics modelling [30]. In the Frozen Vortex Wake method the filaments follow a 'stationary' path away from the turbine. The Free Vortex Wake method convects the filaments based on local velocity, including wake self-induced velocity. The latter one increases model fidelity and is therefore mostly used in the field. An example for the development of such a model is presented by [9]. A variation of this method is the Vortex Particle method where filaments are substituted by vortex particles [31].

2.3.3. Actuator-Based Methods

Actuator-based methods include three main methods with increasing fidelity. First of all, the Actuator Disk Method (ADM) [32] uses Blade Element theory to calculate the aerodynamic loads inside annuli. The undisturbed inflow velocity is originates from the Navier-Stokes Equations (NSE) implemented in CFD. The total load of N blade sections is distributed over the cells in each annulus. The force in a cell is applied as a source term in the NSE. This forcing affects the flow towards the next time step. This method implies a infinite number of blades resulting in an inaccurate near-wake modelling. Secondly, the ALM [33] is a response to this observation. In this method the blades consist of an array of actuator points where the velocity is sampled and forces are projected in the flow. More detail on this method is provided in Section 4.3. Lastly, the Actuator Surface Method (ASM) [34, 35] includes the chordwise load distribution of the blade to add accuracy to the near wake. The challenge for this method is the need for pressure distribution data of the airfoils of the blade.

Since these methods rely on two-dimensional airfoil data and do not resolve the blade surface, they inherently miss three-dimensional flow phenomena such as stall delay and dynamic stall are not included. For this reason various correction methods exist for the polars of the airfoils. Lastly, the methods do show convergence issues due to inaccurate sampling of the velocity and due to suboptimal force projections [28].

2.3.4. Blade-Resolved CFD

Blade-resolved CFD is the most computationally expensive of the modelling techniques. In this approach, the full blade geometry is explicitly meshed, and a no-slip boundary condition is applied on the blade surface to capture the boundary layer. To accurately resolve this layer, the mesh near the wall must be sufficiently refined to represent the shear at the surface. This high resolution significantly increases computational cost.

The main advantage of this method is that it does not rely on two-dimensional airfoil polars to model the blade geometry. As a result, three-dimensional flow phenomena such as tip loss, stall delay, and dynamic stall are inherently captured.

2.4. Experimental Methods for Testing FOWTs

Experimental testing of wind turbines is essential for validating numerical models. Ideally, experiments would be conducted on full-scale turbines operating under realistic atmospheric and oceanic conditions. However, this approach presents significant challenges: it is costly, the flow field is difficult to visualise due to the large spatial scales, and the inflow conditions cannot be controlled. As a result, most tests are performed in wind tunnels, where the flow is known and controllable. For FOWTs, wave basins or motion platforms are incorporated to simulate the dynamic behaviour induced by wind and waves.

Despite the challenges, some full-scale flow measurements have been carried out [36]. The largest wind turbine tested in a wind tunnel to date is the NREL Phase VI turbine[37], which was evaluated in the NASA Ames wind tunnel with a cross-section of 24.4 m by 36.6 m. This campaign primarily focused on pressure distribution over the blade sections, with supplementary smoke visualisation to observe helical tip vortices.

More recently, [38] performed super-large-scale on-site Particle Image Velocimetry (PIV) on a 2.5 MW turbine. This method used natural snowflakes illuminated by laser light to visualise the flow field both upstream and downstream of the rotor, allowing velocity fields to be reconstructed from image sequences.

For experimental testing of floating wind turbines in controlled environments, two primary approaches are available [39]: wave basins equipped with one or more fans to generate wind, and wind tunnels incorporating a motion platform to simulate platform dynamics. Since neither setup can accommodate a full-scale wind turbine, scaled-down models must be used, which introduces additional challenges related to scaling laws and dynamic similarity.

2.4.1. Wave Basin

A wave basin typically consists of a wave maker, a test area, and a wave absorption zone [39]. The wave maker imparts energy to the water to generate waves, which then propagate toward the test area where the floating wind turbine model is located. To prevent reflections that could interfere with the main wave, the downstream end includes an absorption zone designed to dissipate the wave energy.

Wind is usually generated using one or more fans. However, unlike wind tunnels, wave basins lack flow-conditioning devices and generally have rough, non-aerodynamically optimised walls. As a result, the wind conditions are typically less uniform and more difficult to control. Therefore, wave basin testing is best suited for studying the dynamic response of the floater under combined wind and wave loading, while detailed flow measurements around the turbine and floater tend to be less reliable.

A consequence of putting the emphasis on the floater dynamics rather than the aerodynamics is that Froude scaling must be adopted for downscaling the structure. This scaling is discussed in Section 2.4.3.

2.4.2. Wind Tunnel

Wind tunnels offer well-controlled flow conditions but introduce challenges in reproducing realistic platform motions. Typically, motion is imposed at the turbine base via one or more mechanical joints, allowing between one and six DOFs.

[40] employed PIV to measure the wake of a 1:400 scale turbine model equipped with a joint enabling pitch motion. A tunable weight below the joint allowed for adjustment of the motion frequency. However, the setup featured a mean pitch angle of 17.6° , a very low reduced frequency, and a low normalised motion velocity—conditions that are not representative of real FOWT dynamics. In a follow-up study, [41] placed a second 1:400 model in the wake of the first, also in pitch motion, and again used PIV to analyse the resulting flow.

[42] implemented surge motion using a rail-mounted turbine and also utilised PIV for wake measurements, but this setup similarly suffered from non-representative motion frequencies. In contrast, [43] employed a motion platform capable of both surge and pitch, allowing for more realistic motion dynamics. This setup was later reused in [44] to measure wake velocities using a traversing anemometer.

A more advanced approach involves using a motion platform capable of all six DOFs, known as a hexapod. Extensive research at Politecnico di Milano has tested a scale-model FOWT mounted on such a robotic substructure. This configuration supports hybrid testing, as discussed in Section 2.4.4. Key studies using this setup include [45, 46, 47], with further investigations conducted in [48, 49, 13, 14].

2.4.3. FOWT Scaling

FOWT scaling is based on similarity principles, which aim to preserve the dynamic behaviour of the scaled model by ensuring that relevant forces scale appropriately. A wind turbine is subjected to both gravitational and aerodynamic loads, which scale differently—leading to what is often referred to as a scaling effect or scaling mismatch.

Gravitational loads and the stability of the FOWT are governed by the Froude number, defined as the ratio of inertial to gravitational forces:

$$Fr = \frac{U}{\sqrt{Lg}} \quad (2.7)$$

If the model is geometrically scaled while maintaining a constant Froude number, a length scaling factor λ_L is introduced, defined as the ratio of full-scale to model-scale characteristic lengths. From this, the velocity scaling factor can be derived:

$$\lambda_U = \sqrt{\lambda_L} \quad (2.8)$$

In contrast, aerodynamic forces depend on the Reynolds number, the ratio of inertial to viscous forces, given by:

$$Re = \frac{\rho U L}{\nu} \quad (2.9)$$

In practice, the choice of scaling often depends on experimental constraints. For example, in wave basins, where wind quality tends to be poor but floater motion can be accurately captured, Froude scaling is typically preferred to ensure correct hydrodynamic response. Conversely, in wind tunnel experiments where the platform motion is externally imposed and not a result of hydrodynamic forces, Reynolds scaling may be more appropriate to better capture aerodynamic behaviour.

A more pragmatic alternative is performance-based scaling, where the turbine is scaled to match key aerodynamic performance metrics, such as thrust and power coefficients. For instance, [50] scaled the freestream velocity so that the maximum normalised pitch-induced velocity matched between the model and full-scale turbine.

2.4.4. Hybrid/Hardware-In-the-Loop Testing

A way to overcome the scaling mismatch is the hybrid/Hardware-In-The-Loop (HIL) methodology, explained extensively by [45]. This approach couples numerical simulation and physical testing by integrating the Equations of Motion (EOMs) of the FOWT, as shown in Equation 2.10:

$$[\mathbf{M}]_f \ddot{\mathbf{x}} + [\mathbf{K}]_f \mathbf{x} = \mathbf{F}_{floater} + \mathbf{F}_{aero} \quad (2.10)$$

Here, \mathbf{x} is the turbine translational and rotational displacement. The left hand side contains the numerically calculated/approximated mass and stiffness of the the floater with the turbine displacement/rotation in all six DOFs. The first right hand side term is the numerically calculated hydrostatic, mooring and hydrodynamic load due to waves and current. The last term is the aerodynamic load which is obtained from a load cell in the model turbine. The force measured from the load cell does not only contain the aerodynamic load but also inertial effects of the physical turbine. These can be eliminated from the load measurements to isolate the pure aerodynamic load.

By physically reproducing the aerodynamics of the turbine and numerically reproducing the floater dynamics, this method solves the scaling mismatch. It allows for freely scaling the length and velocity scaling factors to fit the wind tunnel characteristics as is done in [51].

2.4.5. Reynolds Number Mismatch

One of the earliest investigations was conducted by [36], who showed that the wake centreline velocity ratios of a model turbine agreed well with those of a full-scale turbine. [52] used PIV to measure the wake across a range of $Re_c = 6400 - 16000$ and tip speed ratios. They concluded that the near wake shared key characteristics with that of a full-scale turbine and that the wake was nearly Reynolds-number independent when the TSR was kept constant.

[53] tested a scaled wind turbine across $Re_D = 1.66 \times 10^4 - 1.73 \times 10^5$ and observed that, for diameter-based Reynolds numbers above 9.3×10^4 , key flow statistics—such as the mean velocity, TI, Reynolds stress, and velocity skewness—became independent of Reynolds number.

Similarly, [54] conducted a series of experiments on a model turbine over a range of $Re_c = 3620 - 31400$, showing that at low Reynolds numbers, both the thrust coefficient and wake expansion deviated significantly from higher Reynolds number cases. They recommended redesigning rotors specifically for low-Reynolds-number conditions.

This approach is adopted in more recent work, such as [51], where low-Reynolds-number airfoils are used to match the thrust curve of a full-scale turbine.

2.5. Numerical–Experimental Evaluation of FOWT Aerodynamics

FOWTs experience platform motions that significantly influence their aerodynamic behaviour. This section reviews the progression of research on FOWT aerodynamics, from early simulations to recent experiments, with a focus on the effect of surge and pitch motion, less-studied motion DOFs, and combined system responses including blade flexibility and control. The aim is to highlight key findings, identify recurring patterns, and expose remaining knowledge gaps.

2.5.1. Initial research

The first aerodynamic studies on FOWTs were conducted by [9], who developed a free-wake vortex method to simulate operating conditions specific to floating platforms. Their results revealed an increased likelihood of unsteady aerodynamic phenomena and highlighted violations of widely accepted assumptions used in conventional wind turbine models, such as the slipstream assumption in BEMT, which implies no reversed flow through the rotor plane.

In a follow-up study, [10] analysed the induction at the rotor caused by wake vorticity. They observed that platform motion generates distinct bunches of shed vorticity in the wake. The resulting induction at the rotor decays rapidly; dropping to just 1% at two rotor diameters downstream. Furthermore, the wake stability was shown to deteriorate under lower inflow speeds or higher TSRs, due to low vorticity convection leading to strong wake self-induction; e.g., via mutual-inductance instability.

In another study, [2] examined the angle of attack variation resulting from floater motion. They found that inboard blade sections exhibited the highest frequency of stall and unsteady effects. Additionally,

they demonstrated that the slipstream assumption becomes invalid for FOWTs, especially at the outboard blade sections under below-rated wind speeds. As a result, classical BEMT models cannot be directly applied to FOWTs.

However, subsequent work by [55] and [27] showed that dynamic inflow models could address these limitations. This was experimentally validated by [56], confirming the potential of dynamic models to more accurately represent FOWT aerodynamics.

2.5.2. Surge and Pitch Motion

Following the initial studies by [2, 9, 10], further research into FOWT aerodynamics has focused particularly on surge and pitch motions, as these have the most pronounced impact on rotor performance.

Recovery

[40] conducted wind tunnel experiments using PIV to study the mean wake statistics under dynamic pitch motion. Their results revealed an upward shift in streamwise kinetic energy and TKE flux, which is not captured in standard engineering wake models, leading to inaccurate wake modelling and longer wake recovery distances. However, as noted in Section 2.4.2, their setup featured a large mean pitch angle (17.6°), a low reduced frequency, and a low normalised motion velocity, limiting its representativeness. The upward shift is later also observed by [57], although in this case faster wake recovery occurred due to wake stretching.

Similarly, [42] performed PIV-based experiments on a surging turbine and observed reduced Reynolds stresses. However, this early experiment also employed a low reduced frequency. In contrast, [58] observed that rotor-wake interaction due to surge motion results in increased kinetic energy at the rotor and elevated TKE at the root. Moreover, they found that surge motion reduces wake recovery, attributed to the high stability of tip vortices. [57] further showed that, under low turbulent inflow, low-frequency motion enhances wake recovery more than high frequencies and is less sensitive to amplitude. Motion with high frequencies were even observed to make the wake more stable, making the recovery slower than for a bottom-fixed turbine. This effect was more pronounced in surge than in pitch motion.

The influence of inflow turbulence was investigated by [59], who demonstrated that turbulent inflow accelerates wake recovery by destabilising the wake. This results in significantly different streamwise velocity profiles at identical downwind positions.

Performance

Regarding rotor performance, [60] performed numerical studies confirming experimental observations by [56], particularly for high TSRs. They found that the increase in power and thrust amplitude with surge is caused by near-wake dynamics: vorticity accumulation leads to greater wake expansion, inducing axial and tangential flow fluctuations at the rotor. Consequently, low-fidelity models such as BEMT become unreliable at high TSRs. Additionally, [61] noted a decrease in power at higher surge frequencies for low TSRs. Experimental work by [48] and [14] showed a quasi-steady relation between surge and pitch motion velocity, thrust and power. Beyond a certain frequency threshold, the thrust amplitude was found to increase more rapidly with frequency, indicating the onset of unsteady aerodynamic effects.

The effect of surge motion on power and thrust dynamics was further examined by [62] using blade-resolved unsteady Reynolds Averaged Navier-Stokes (RANS) simulations. Their study revealed that surge motion amplifies power more than thrust and introduces a -90° phase shift in thrust due to aerodynamic damping. They also observed vortex pairing in the wake, induced by varying vortex shedding velocities throughout the motion cycle. The imposed motion frequency, which was a harmonic of the rotor azimuthal frequency, led to a stable vortex shedding pattern.

[63] adds that the mean thrust decreases due to surge motion, but that the mean power increases with more than 3%. An even higher increase of power is presented by [64], both for pitch and surge motion. This was attributed to the asymmetry in velocity during the motion cycle: the upwind portion contributes significantly to power due to the cubic dependency on velocity, while the downwind portion reduces

power less severely. [58] added that the power gain shows an increasing trend with both amplitude and frequency.

Wake States

The dynamics of wake states during surge and pitch motion were further investigated by [65], who studied the propeller state and VRS using numerical simulations. The VRS occurs when the blade tip interacts with a tip vortex in the wake. For this to happen, a zero or negative relative velocity must be maintained, allowing the blade tip to reach the tip vortex. During this condition, root vorticity also accumulates just downstream of the hub. [64] observed similar phenomena for high surge and pitch amplitudes. In light of the predicted occurrence of the VRS, [66] analysed three different prediction criteria. They concluded that while the Wolkovitch criterion is most effective at predicting the VRS, the Peters criterion is better suited for identifying early aerodynamic changes that may signal its onset. Additionally, [57] found that pitch motion induces the most significant rotor-wake interactions, particularly at higher frequencies. [67] further noted that fast surge motion can lead to a decrease in average power due to stall and the occurrence of the VRS. [68] explored the transition of wake states under surge motion using blade-resolved CFD. They found that before entering the propeller state, the turbine passes through an intermediate propeller break state. They also described a new wake regime—the quasi-windmill state—which exhibits negative thrust but positive power.

Stability

The influence of motion on wake stability and spectrum has also been examined. [69] showed that pitch motion leaves a clear signature in the wake spectrum, reducing energy content below the pitch frequency, though this difference diminishes farther downstream. [48] observed similar effects for surge motion in experimental settings and further showed that surge motion causes variable tip vortex travel speeds. More specifically, [70] investigated tip vortex stability in heave, surge, and pitch motions using linear vortex stability theory. Their study, conducted under low turbulence conditions, identified two forms of instability: the leapfrogging instability and a six-vortex interaction instability occurring at specific frequencies inherent in the tip vortex. The latter leads to strong oscillations in streamwise velocity and may be exacerbated by turbine motion.

Finally, [71] applied low-dimensional non-linear stability theory to show that specific motion frequencies—both in fore-aft and side-to-side directions—can enhance wake recovery. For fore-aft motion, wake pulsation was observed in reduced frequencies between 0.2 and 0.55. In the range from 0.55 to 0.9, a new quasi-periodic coherent structure emerged.

2.5.3. Knowledge Gap

In literature there is no consensus on the effect of the floater motion on the recovery of the wake. However, for FOWT farms its effects are required to optimise for energy capture and life duration of the turbine. This knowledge gap is exploited in this work by analysing experimental results.

Next to this, numerical models are often not validated directly against experimental results since these are unavailable. However, experimental results from [6] allow for a tuning and validation in this work.

Part II

Methodology

Experimental Setup

Numerical simulations of a scale model FOWT are compared to experimental velocity field measurements using the same model. For a meaningful comparison, it is important that both setups are nearly identical. Therefore, this chapter first describes the test environment in Section 3.1. Next, the wind turbine scale model is introduced in Section 3.2. The motion conditions applied to the model are explained in Section 3.3. Finally, the method used to obtain velocity measurements in the wake using PTV is discussed in Section 3.4.

3.1. Test Environment

The experimental campaign was conducted in the Open Jet Facility at Delft University of Technology¹. This wind tunnel includes a 2.85-meter by 2.85-meter nozzle that exhausts a jet into an open test section, which spans 13 meters in width and 8 meters in height.

The jet provides a low-turbulence, uniform inflow with a turbulence intensity (TI) of 0.5% at a distance of 1 meter from the nozzle exit. A detailed flow quality analysis has been carried out by [21].

The temperature is maintained at 20°. The flow exiting the nozzle has a top-hat profile. As it propagates, a shear layer forms between the jet and the surrounding still air, causing the flow to diffuse. This diffusion leads to a freestream jet contraction of approximately 4.75°. Additionally, within the first 4 meters downstream of the nozzle, the TI increases to 1.7%. Radially, the turbulence intensity rises to a maximum of 15% around 1.5 meters from the jet axis.

3.2. Wind Turbine Scale Model

Within the wind tunnel, a 1:148 scale model of the DTU 10 MW wind turbine [11] is installed. The rotor was specifically designed for this scale model, as detailed in [51].

One of the primary design objectives was to replicate the thrust coefficient distribution of the full-scale turbine under below-rated operating conditions. This is critical to achieving a comparable aerodynamic response. Furthermore, since the model is intended for HIL testing with the floater simulated numerically, Froude scaling is not required (see Section 2.4.3). Instead, the model scale is based on the size constraints of the wind tunnel, meaning Reynolds number similarity is not preserved.

The full-scale turbine operates at Reynolds numbers between 1.5×10^6 and 1.4×10^7 , while the scale model operates around 1×10^5 . As explained in Section 2.4.5, a diameter-based Reynolds number exceeding 9.3×10^4 is necessary to capture key wake features. The scale model meets this criterion.

To address the lower Reynolds number, the SD7032 low-speed airfoil is used in place of the original profiles. The airfoil polars were extrapolated for post-stall behavior using the Viterna method [17], and corrected for stall-delay effects via the DuSelig method [16]. These corrected polars guided the optimization of twist and chord distributions across 39 blade sections to match the thrust response of the full-scale DTU 10 MW turbine.

A constant TSR of 7.5 is maintained by setting a fixed rotor speed under below-rated conditions. In contrast to full-scale turbines, the scale model does not include variable pitch control for above-rated conditions. Additionally, the turbine structure is designed to be fully rigid, minimizing any aerodynamic uncertainties due to structural deformation.

¹<https://www.tudelft.nl/lr/organisatie/afdelingen/flow-physics-and-technology/facilities/low-speed-wind-tunnels/open-jet-facility>

The final geometry and operating parameters are summarized in Table 3.1. The rotor diameter is 1.2068 meters. The tower is a rigid cylinder measuring 0.81 meters in height with a diameter of 8 cm. The hub is positioned 0.037 meters above the tower top, and an overhang of 0.095 meters provides clearance between the rotor and tower, eliminating the need for shaft tilt. The rated wind speed of 4 m/s results in a rated rotor speed of 480 rpm.

Table 3.1: Wind turbine scale model design parameters.

Parameter	Value	Unit
Diameter	1.2068	m
Tower height	0.81	m
Tower-to-shaft	0.037	m
Shaft tilt	0	°
Overhang	0.095	m
Rated wind speed	4	m/s
Design TSR	7.5	

3.3. Motion Cases

The turbine is mounted on a six DOFs parallel robotic motion platform². In this study, only surge and motion are examined.

As discussed in Section 2.5, it was found that a rotor-based reduced frequencies below 1.2, corresponding to a frequency of 4 Hz for the scale-model turbine, resulted in a quasi-steady thrust response [13]. This means that the thrust amplitude is linearly proportional to the motion induced normalised velocity ΔV^* :

$$\Delta V^* = \frac{V_{hub}}{U_0} \quad (3.1)$$

where V_{hub} is the velocity of the hub and U_0 is the normalisation velocity, chosen to be equal to the inflow velocity.

Any value of ΔV^* can be chosen. However, a value of 0.1 has been chosen since this maximises the hub motion and thus the flow response, while frequencies representative for waves could be analysed without reaching the actuation limits of the motion platform.

Two frequencies of 2 Hz and 5 Hz have been chosen, which are equivalent to rotor diameter based reduced frequency of 0.6 and 1.5, respectively. In this report they are referred to as "low-frequency" and "high-frequency" cases. The former is representative for the rigid body natural frequency of the turbine. The latter represents wave frequencies. These frequencies have been chosen to understand the differences in the wake of a turbine with quasi-steady and unsteady motion. The high-frequency case is limited by the motion platform capabilities. The exact test matrix is presented in Table 3.2

Table 3.2: Motion cases test matrix.

f [Hz]	ΔV^*	A_s [mm]	A_p [°]
2	0.1	31.8	1.69
5	0.1	12.7	0.68

3.4. Particle Tracking Velocimetry

To measure the velocity field in the wake of the wind turbine, the state-of-the-art "Shake-the-Box" 3D Lagrangian PTV setup, developed at Delft University of Technology, is employed. The full experimental setup is shown in Figure 3.1, and is further detailed in [6].

²<https://www.quanser.com/products/hexapod/>

excluded due to significant flow deterioration in this region, attributed to boundary roughness effects, mixing of the jet with the surrounding air and flow blockage due to the raster in the outflow, which introduced excessive uncertainty.

Another notable uncertainty relates to the buoyancy of the helium-soap mixture. During analysis, it was observed that the wake appeared to shift downward with increasing downstream distance. This could be due to a small miscalibration in the camera alignment, due to a slightly negatively buoyant helium-soap mixture, or due to a low pressure region behind the table where the turbine was placed on.

Although this manual alignment was corrected by visual inspection, its precision is limited. Comparing the experimental data with numerical simulation results, as done in this work, allows for a more accurate registration of volumes by applying corrections (see Appendix D), demonstrating a valuable synergy between experimental and computational approaches.

Numerical Setup

The experimental setup (see Chapter 3) is replicated in AspFAST. This is a coupling between the LES code GRASP and the multi-physics, multi-fidelity tool OpenFAST. The structure of this coupling is shown in Figure 4.1. The AspFAST module receives inputs from the application programming interfaces of both GRASP and OpenFAST.

In this chapter, the responsibilities of the different modules are discussed. Section 4.1 covers the LES GRASP, followed by a discussion of OpenFAST in Section 4.2. Subsequently, the AspFAST module is addressed in Section 4.3. Section 4.4 discusses the used computational resources for this work. Lastly, the method for assessing the simulation results is explained in Section 4.5.

Although the modules are extensive and offer many configuration options, only the relevant options will be discussed in this chapter. Furthermore, not the entire setup will be defined here; the initial setup will be presented in Chapter 5.

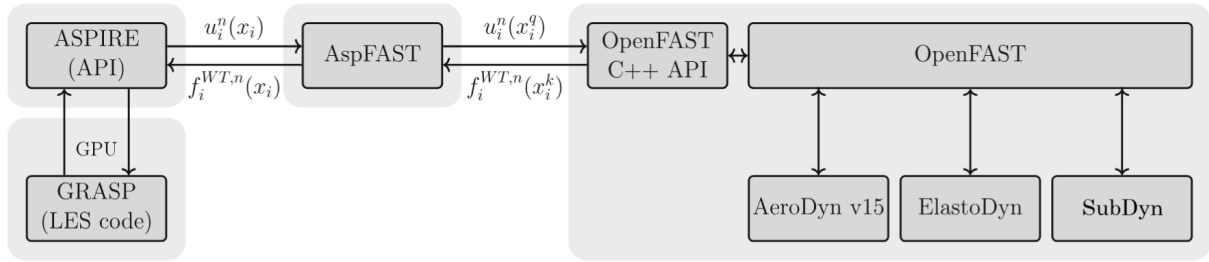


Figure 4.1: Schematic of the numerical framework (modified from [7]). The central coupling module AspFAST integrates the LES solver GRASP, developed by Whiffle, via the ASPIRE interface (left), with OpenFAST (right). For this study, the Aerodynamics, Elastodynamics, and Sub-structure dynamics modules of OpenFAST are used.

4.1. GRASP: Large Eddy Simulation

GRASP [72] is a CFD LES code developed by Whiffle Weather Finecasting Ltd¹. It is an atmospheric LES model designed for high-precision weather forecasting in the renewable energy sector. Thanks to its implementation on a highly parallel Graphical Processing Units (GPU) architecture, it operates significantly faster than conventional Central Processing Units (CPU)-based counterparts.

4.1.1. Computational Fluid Dynamics

Conventional CFD methods typically employ the finite-volume approach, in which the computational domain is discretised into control volumes (cells) and time steps. The NSE are then solved numerically on each cell at each time step using integration and interpolation techniques. The equations for incompressible flow are given in Equation 4.1. In wind energy applications, the flow is commonly assumed to be incompressible, as velocities remain below Mach 0.3. Although recent studies have questioned this assumption for large-scale turbines [73], it remains valid for the scale model considered here, which has a tip speed of 30 m/s.

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} &= -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned} \tag{4.1}$$

¹<https://whiffle.nl/>

The flow around a wind turbine is turbulent, and the nature of turbulent flow is commonly described using the concept of the energy cascade. Turbulent flows contain a wide range of structures with varying spatial and temporal scales. The large scale structures contain much energy and have low frequency/wave number, while the small scales have little energy and a high frequency/wave number due to the low time and spatial scale. The high energy at low frequency/wave number is called the energy-containing range. Here, the mean flow feeds large turbulent structures with kinetic energy. During the inertial subrange the structures break down into smaller structures. At some point the smallest possible scales are reached inside the dissipative range. These structures are called the Kolmogorov scales and mainly dissipate kinetic energy into heat [25].

To solve the NSE exactly, all turbulent structures must be resolved in both space and time, including the smallest scales. This requires discretising the computational domain and time steps with resolutions finer than the Kolmogorov scales.

Simulations that achieve this level of resolution are known as direct numerical simulations (DNS). However, the computational cost of DNS increases with the Reynolds number to the power of 4.5, as the smallest scales become increasingly fine [25]. For high-Reynolds-number flows such as those encountered around wind turbines, DNS is computationally infeasible.

Large Eddy Simulation

LES provides a more efficient alternative by explicitly resolving only the large turbulent structures while modelling the smaller, SGS motions. This approach effectively applies a spectral filter to the energy cascade, removing high-frequency components.

The unresolved scales are collectively described by a turbulence SGS model. Among these, eddy-viscosity models are the most widely used. These models add artificial (eddy) viscosity to the physical viscosity with approximate the dissipative effects of the unresolved scales.

The added viscosity depends on a model dependent coefficient C_x , a characteristic filter width Δ (related to the cells size) and operator $\bar{\bar{\omega}}$:

$$\nu_{SGS} = (C_x \Delta)^2 \bar{\bar{\omega}} \quad (4.2)$$

4.1.2. General Inputs

GRASP requires a number of key settings. These include the total runtime and maximal time step. The runtime is chosen such that measurement time is equal to 10 seconds, equivalent to the experiments. A transient phase occurs during wake development, the duration of which depends on the domain length. It has been found that the freestream flow should pass through the downwind portion of the domain approximately three times to capture a developed wake.

Time integration is performed using a third-order Runge-Kutta scheme. The time step is set small enough to ensure numerical stability and accuracy. The flow velocity is limited such that the CFL number remains at or below 1, meaning the flow can travel at most one cell per time step. However, as will be discussed in Section 4.3, an even smaller time step is required to maintain the stability of the ALM.

Momentum advection schemes determine how velocity is spatially integrated. These schemes advance the velocity field in space and differ in accuracy and computational cost. Generally, higher-order schemes provide greater accuracy but are more computationally expensive. GRASP offers a second-order central difference scheme and a fifth-order scheme. The second-order scheme is chosen here due to its widespread use in the literature [7]. However, as can be observed in Appendix B a higher order scheme can improve the results.

4.1.3. Domain Inputs

In GRASP the CFD mesh must be defined. A number of cells along the length, width and height must be set. Beside that, the domain size is set in metres which implies a certain cell size.

Mesh refinement is only possible in the vertical (z) direction towards the bottom boundary. This limitation arises because GRASP was originally designed to model the atmospheric boundary layer, where high resolution near the surface is necessary to resolve turbulent structures [25]. Unfortunately, this vertical-only refinement is not ideal for the present work since no atmospheric boundary layer is modelled. Instead, mesh refinement in the rotor wake region, commonly used in literature [59], would be more beneficial to reduce numerical diffusion of vorticity; especially around the helical tip vortices.

For performance optimisation on GPU architectures, it is recommended that the total number of cells be a multiple of 32. This is due to the GPU thread organisation in 'warps' of 32 threads. If the number of cells is not a multiple of 32, some threads remain idle, resulting in inefficient resource use.

The domain length is chosen to provide sufficient upwind and downwind distances for accurate wake analysis and comparison with experiments. Generally, the turbine should be positioned far enough from domain boundaries to avoid spurious interactions that could degrade accuracy. The required domain dimensions will be discussed further in Chapter 5.

4.1.4. Initial Conditions

GRASP requires initial conditions for variables such as velocity, temperature, and humidity in all mesh cells. For this work, a uniform inflow is specified. Two vertical levels are sufficient, each with a constant velocity in the x -direction ($u=4$ m/s). Additionally, a constant temperature of 293 K and a pressure of 1×10^5 Pa are prescribed for convenience.

Restart

GRASP supports restarting from a previously saved flow field, which improves robustness in case of interruptions. This is particularly useful when running simulations on computing clusters that are subject to execution time limits or unexpected errors.

The restart field must include, at a minimum, the three velocity components, temperature, humidity, and precipitation. It must also match the dimensions and resolution specified in the configuration of the new simulation.

4.1.5. Boundary Conditions

The boundary conditions are divided over the bottom boundary condition and the boundary conditions for the other sides. The boundary condition at the base-level is a no-slip wall condition and the top boundary condition is a slip wall condition. All the other boundaries are periodic. To understand the boundary conditions and the way to manipulate them, you must understand the concept of 'nudging'.

Nudging

Nudging refers to the coupling of two parallel simulations via relaxation at their boundaries. These are typically referred to as the precursor and cursor simulations. The precursor simulation develops a particular flow field in a domain of the same size as the cursor. A nudging extent is specified within the cursor domain, covering a number of cells where a gradual relaxation from one flow to the other occurs. This ensures consistency between the two simulations at their interface.

The nudging operation for a scalar s is defined by Equation 4.3, where subscripts 1 and 0 refer to the cursor and precursor simulations, respectively. The relaxation strength $\alpha_n(x, y, z)/\Delta t$ depends on the extent and can be scaled using the nudging factor η_n .

$$s_1(x, y, z) = -\frac{\alpha_n(x, y, z)\eta_n}{\Delta t} [(s_1(x, y, z) - s_0(x, y, z))] \quad (4.3)$$

Inflow and Outflow Boundaries

The inflow and outflow boundaries are periodic. This means the wake leaving the domain at the outflow re-enters through the inflow boundary, leading to a feedback loop. To prevent this, nudging is applied

from the precursor to the cursor simulation, enforcing a uniform inflow. Since the precursor does not contain the wind turbine, its flow remains undisturbed.

Choosing appropriate nudging parameters is non-trivial. The nudging extent reduces the physically accurate domain, so numerical efficiency requires it to be minimal. However, if the extent or factor is too small, the inflow may become non-uniform. On the other hand, setting the factor too high (above 1) can lead to instabilities due to overshooting. In Chapter 5, the minimum required nudging extent for a factor of 1 is determined.

Lateral Boundaries

The lateral boundaries are also periodic and subject to nudging from the cursor to the precursor simulation. This effectively imposes a velocity at the boundaries that may interfere with the turbine wake. This phenomenon resembles the blockage effect, in which lateral boundaries impose pressure on the flow, leading to reduced wake expansion and increased velocity through the rotor disk due to the Venturi effect.

According to [31], blockage effects are mitigated when the blockage ratio is kept below 10%, and ideally below 5%. Therefore, to reduce the influence of the lateral boundaries, the cross-sectional area of the domain should be sufficiently large, as discussed in Chapter 5. If the domain is wide enough, the nudging regions at the lateral boundaries can be kept small without introducing significant distortions.

Bottom and Top Boundaries

The bottom wall has a no-slip boundary condition, which induces shear stress and causes the flow to develop a logarithmic profile. Such a profile is undesirable in this work, as it conflicts with the uniform inflow conditions of the wind tunnel. Fortunately, the shear can be suppressed in the NSE through a configuration in GRASP. Additionally, the roughness length z_0 must be smaller than 0.1 to avoid numerical instabilities.

Because the bottom and top boundaries are not periodic, the top boundary is defined as a slip wall. This boundary condition is described in Equation 4.4, where the vertical velocity is zero and vertical gradients of horizontal velocities are also zero. Scalar quantities follow a constant vertical gradient, initialised at the start of the simulation.

$$\begin{aligned} \frac{\partial u}{\partial z} = \frac{\partial v}{\partial z} = w = 0 \\ \frac{\partial s}{\partial z} = c_s \end{aligned} \tag{4.4}$$

Turbulence

Turbulent inflow may be desired. GRASP does not provide direct control over TI or TKE, but instead relies on natural processes to introduce flow unsteadiness.

In nature, turbulence arises from shear stresses, surface-induced temperature and humidity variations. Solar heating causes spatial differences in air temperature and density, while moisture variation results from uneven evaporation. These differences generate buoyancy-driven fluctuations in the flow.

GRASP replicates these processes through user-defined vertical heat and/or humidity fluxes. Additional parameters include the random seed, and the magnitude of temperature and humidity perturbations.

Although heat and humidity fluctuations are negligible in a wind tunnel, they can still be used to generate the desired turbulence. Only one (heat or humidity) is needed to initiate turbulence, as both affect air density similarly. To minimise aerodynamic effects around the turbine, the flux is kept small (order of 1×10^{-1} K/s). The relation between flux, perturbation magnitude, and resulting TI is further analysed in Chapter 5.

Since turbulence takes time to develop, it should be generated in a separate precursor simulation before the cursor simulation begins.

4.1.6. Subgrid-Scale Model

GRASP contains various SGS models: Smagorinsky, Sullivan, Vreman, Verstappen, Rozema and Uniform [74, 75, 76, 77, 78]. Each model has a model constant C_x and a filter width Δ which is used to calculate the SGS eddy viscosity. The Uniform SGS model simply specifies a constant eddy viscosity throughout the complete domain. In the inputs this is specified as K_m and no model constant is needed.

The Smagorinsky model is used as a baseline since it is widely adopted in the literature due to its simplicity [13, 59, 79]. A value of 0.17 is used for the model constant, which is for isotropic turbulence [25]. However, the Verstappen and Rozema models may offer improvements for modelling anisotropic turbulence, which the Smagorinsky model is known to represent poorly.

Smagorinsky Model

The Smagorinsky model [74] was the first SGS model to be developed [25]. Its eddy viscosity operator depends solely on the resolved strain-rate tensor \tilde{S}_{ij} .

$$\nu_{sgs} = (\Delta C_x)^2 \sqrt{2\tilde{S}_{ij}\tilde{S}_{ij}} \quad (4.5)$$

The model performs well in constant flow regimes, but when transitions between regimes occur, the model constant must be adapted [25]. To address this, the Dynamic Smagorinsky model has been proposed, where the model constant adapts to the local flow. While this allows for transitional flow modelling, it can result in unphysical negative values for the constant.

Verstappen and Rozema Models

The Verstappen model [77] is a minimum-dissipation model, designed to provide only the necessary amount of turbulent dissipation. It relies on the trace of the second and third invariants of the strain-rate tensor. Its formulation is given in Equation 4.6. The model has several useful properties: it yields zero viscosity in laminar regions, switches off for backscatter or two-dimensional flow, and produces dissipation consistent with the exact SGS tensor on isotropic grids [78].

$$\nu_{sgs} = (\Delta C_x)^2 \frac{\max(\int_{\Omega_\Delta} r(u)dx, 0)}{\int_{\Omega_\Delta} q(u)dx} \quad (4.6)$$

where

$$r = -\frac{1}{3}tr(S^3) \quad (4.7)$$

$$q = \frac{1}{2}tr(S^2) \quad (4.8)$$

However, this model's performance depends on the definition of the filter width. As a result, its dissipation is not consistent on anisotropic grids. The Rozema model [78] addresses this limitation by modifying the operator used in the Verstappen model, resulting in improved performance.

In the original Verstappen model, an optimal model constant of 0.289 was reported. Later work [80] suggested a value of 0.356. Rozema et al. [78] found that these values were still too low and depended on the numerical scheme. For a second-order scheme, they reported $C_x = 0.333$, and for a fourth-order scheme, $C_x = 0.486$. The Rozema model has slightly different optimal constants: 0.5480 for a second-order scheme and 0.460 for a fourth-order scheme.

GRASP allows mesh refinement only in the vertical direction, which is aimed at resolving the atmospheric boundary layer. Since this work does not include such boundary layer effects, an isotropic grid is used. Under these conditions, the choice between the Verstappen and Rozema models becomes arbitrary, provided a suitable model constant is selected.

4.2. OpenFAST: Wind Turbine Characterisation

OpenFAST² is a multi-physics, multi-fidelity tool for simulating the coupled dynamic response of wind turbines. It couples various modules ranging from the turbine architecture to wind and wave conditions acting on the turbine.

For this work, the ElastoDyn, AeroDyn and SubDyn modules are used for storing and calculating the state of the turbine. These modules are related to the structural dynamics (Elasto), aerodynamics (Aero) and the sub-structural dynamics (Sub), respectively, and will be further discussed in the following sections.

Each module provides various output options³. It is important to use the correct OpenFAST version (v3.4.0 in this work).

4.2.1. Structural Dynamics

ElastoDyn handles all structural aspects of the turbine, including the tower, nacelle, and rotor. It is responsible for specifying the geometric configuration of the turbine, such as the tip radius, hub radius, rotor overhang, tower-to-shaft distance, and tower height. The initial rotational speed of the rotor is also defined here and is set to 480 rpm, as described in Chapter 3.

In addition, the mass, stiffness, and damping properties of the blades and tower are specified. However, in this work, the turbine is modelled as fully rigid to replicate the design of the scale-model turbine. As a result, these structural parameters are trivial.

ElastoDyn also defines the floater DOFs, including surge and pitch. These degrees of freedom must be explicitly activated if motion is desired. Initial conditions for displacement, rotation, and velocity can also be set.

4.2.2. Aerodynamics

AeroDyn handles the aerodynamic modelling of the turbine. Most importantly, it includes the definition of airfoil lift and drag polars for various blade elements across different Reynolds numbers, as well as the spanwise position of these elements, and the chord and twist distributions of the blades.

Additionally, the environmental conditions, namely air density and viscosity, are specified. The air density is used in Equation 2.2 and Equation 2.3, while the viscosity is used in the calculation of the Reynolds number (see Equation 2.9). The lift and drag coefficients are interpolated from the provided polars as functions of Reynolds number and angle of attack.

If the tower is included in the simulation, this must be explicitly enabled. The aerodynamic properties of the tower, such as diameter and drag coefficient, must also be defined. In this case, the tower diameter is 8 cm, and the drag coefficient is set to 0.5, corresponding to that of a smooth cylinder [81].

4.2.3. Sub-Structure Dynamics

SubDyn defines the external motion of the turbine sub-structure and is used in this work to impose a prescribed sinusoidal motion on the wind turbine. The sub-structure is modelled as a linear mass–damper–spring system using Guyan/Craig-Bampton reduced-order 6x6 mass, damping, and stiffness matrices. These matrices describe the dynamic behaviour in all six DOFs of the floater. An input file provides these matrices along with a time series of external loads for each DOF, which together define the system’s forcing function.

Forcing

The system is governed by the equations of motion, given in Equation 4.9, where only surge and pitch are active in this work:

$$\mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{C}\dot{\mathbf{x}}(t) + \mathbf{K}\mathbf{x}(t) = \mathbf{F}(t) \quad (4.9)$$

²<https://openfast.readthedocs.io/en/main/>

³https://openfast.readthedocs.io/en/main/_downloads/3f19498a5dc774461e022b671ff01ec6/OutListParameters.xlsx

To impose a prescribed sinusoidal motion, the theory of a damped driven harmonic oscillator is applied, as described in Equation 4.10. Here, the mass is set arbitrarily high to suppress sensitivity to noise in the forcing signal. This is acceptable since inertial effects are not relevant in this study:

$$M\ddot{x} + C\dot{x} + Kx = F_0 \cos ft \quad (4.10)$$

The solution to this equation includes both transient and steady-state components [82]. To ensure a steady sinusoidal response, the system is tuned to resonate, which requires:

$$f^2 - f_0^2 = f^2 - \frac{K}{M} = 0 \quad (4.11)$$

The damping is controlled using the damping ratio ζ , defined as the ratio of system damping to critical damping:

$$\zeta = \frac{C}{2\sqrt{KM}} \Rightarrow c = 2M\zeta f_0 \quad (4.12)$$

A damping ratio of $\zeta = 0.7$ is used, which produces an underdamped response with small oscillations and rapid decay of transients. According to [82], the response amplitude near resonance increases rapidly as ζ decreases. Underdamping also avoids overshooting and delays associated with overdamped systems.

With the system fully defined, it is possible to calculate the forcing amplitude F_0 from the solution in Equation 4.13 [82].

$$x = A \cos ft \quad (4.13)$$

$$A = \frac{F_0}{M \sqrt{(f_0^2 - f^2)^2 + \frac{C^2}{M^2} f^2}} \Rightarrow F_0 = ACf \quad (4.14)$$

Pitch motion

OpenFAST requires a turbine with floater motion to have a turbine base at sea level. However, to reduce boundary interference in the CFD simulation, the rotor is positioned near the centre of the domain. As shown in Figure 4.2, this requires extending the tower from its original height $H_{t,0}$ to a new height H_t .

While this has no effect on translational motion, it increases the linear velocity at the rotor during pitch motion due to its greater distance from the centre of rotation. A pitch oscillation with amplitude A_p now induces additional surge and heave motions with amplitudes A_s and A_h , respectively:

Using simple trigonometry, these amplitudes are:

$$\begin{aligned} A_s &= (H_t - H_{t_0}) \sin A_p \approx (H_t - H_{t_0}) A_p \\ A_h &= A_s \tan A_p \approx A_p^2 (H_t - H_{t_0}) \end{aligned} \quad (4.15)$$

To maintain the correct phase, when pitch reaches maximum positive velocity, surge reaches maximum negative velocity and heave reaches maximum positive velocity. The resulting oscillations are:

$$\begin{aligned} x_p &= A_p \sin 2\pi f_p t \\ x_s &= A_p (H_{t_0} - H_t) \sin 2\pi f_p t \\ x_h &= A_p (H_t - H_{t_0}) \tan A_p \sin 2\pi f_p t \end{aligned} \quad (4.16)$$

The maximal pitch angle is only 1.69° , resulting in a maximal heave amplitude that is 3% of the surge amplitude. Therefore, this is considered negligible.




Figure 4.2: Schematic of the pitch motion with amplitude A_p . Due to the extended tower height from H_{t_0} to H_t , the motion must be compensated with a surge and heave oscillation with amplitudes A_s and A_h .

4.3. AspFAST: Actuator Line Method

AspFAST combines the flow field information from GRASP and turbine properties in OpenFAST in an open-source ALM implementation⁴.

The ALM, as developed by [33], discretises the blade into elements with each a blade node, called an actuator point. OpenFAST has all the properties of the element, such as aerodynamic polars and blade geometry, and gets flow inputs from GRASP such that it can calculate a discrete force per unit span \mathbf{f}_{AL} .

AspFAST then takes these discrete forces from OpenFAST and smoothly distributes (or convolves) them over the CFD mesh cells, as presented in Equation 4.17. This smoothing avoids spurious oscillations in the flow field that can result from sharp, discontinuous force application [33]. Each mesh cell receives a body force contribution from all actuator points within its influence region, and these forces are subsequently applied as source terms in the NSE.

$$\mathbf{f}_e = \mathbf{f}_{AL} \otimes \eta \quad (4.17)$$

After the source terms are applied, the NSE are advanced in time and the flow field is updated. This updated flow field is then used to recalculate the forces, and the process repeats.

In the ALM, the time step must be smaller than the time it takes for the blade tip to traverse a single grid cell. In low-speed flows, this constraint is often more restrictive than the Courant Friedrichs Lewy (CFL) condition and is therefore adopted throughout this work.

4.3.1. Body Force Projection

In this work, the force around each actuator point is distributed using the isotropic Gaussian kernel in Equation 4.18, originally proposed by [33]. This formulation concentrates the load in nearby cells while gradually diminishing its influence with distance, following an exponential decay. How far this distribution reaches depends on the kernel size ϵ .

$$\eta = \frac{1}{\epsilon^3 \pi^{3/2}} e^{-\left(\frac{|\mathbf{r}|}{\epsilon}\right)^2} \quad (4.18)$$

⁴<https://gitlab.com/whiffle-public/aspfast/>

According to [83], the number of actuator points per blade should lie between 30 and 60 and should not exceed the number of grid cells. The optimum for this study is examined in Chapter 5.

[33] recommended a kernel size of at least two grid cells to avoid spurious flow oscillations. For this reason, AspFAST defines kernel size relative to the (uniform) grid spacing. Figure 4.3 illustrates the influence of kernel size by showing the summed kernel weights per cell for $\epsilon/\Delta x = 2$ and $\epsilon/\Delta x = 5$. Larger kernel widths effectively increase the blade's apparent thickness and distribute forces over a broader area. The blue dots denote 50 actuator points per blade. The optimal kernel size for AspFAST is evaluated in Chapter 5.

A larger kernel size leads to a more diffuse bound vortex [84]. This also results in diffused tip and root vortices, causing less downwash at the force nodes and thus a performance overprediction.




Figure 4.3: Contours of the summed force distribution kernel η for $\epsilon/\Delta x$ equal to 2 (left) and 5 (right). The blue dots are the actuator points of the actuator lines.

Filtered ALM

The optimal kernel size has been found to be approximately a quarter of the local chord length [85]. To maintain $\epsilon \geq 2\Delta x$, this requires at least eight cells along the chord direction per blade section. In GRASP, local mesh refinement is not available, meaning the entire domain must use this fine resolution, which is computationally expensive.

To address the inaccuracy caused by coarse meshes, [86] developed a lifting line-based correction model. It accounts for the induced velocity deficit resulting from suboptimal kernel widths, particularly in the tip and root regions. The correction treats the smoothed force field as a filtered velocity field and uses an analytical expression for the optimal and actual downwash. The corrected velocity is then given by:

$$\mathbf{U}_\infty = \tilde{\mathbf{u}} + [u_y(\epsilon^{opt}) - u_y(\epsilon^{LES})] \mathbf{j} \quad (4.19)$$

Here, \mathbf{j} is the unit vector perpendicular to the span and lift directions. The downwash velocity $u_y(\epsilon^{opt})$ is obtained from canonical lifting line theory, assuming a Gaussian-filtered bound vortex, constant lift slope, and small spanwise gradients of chord and kernel width. Derivations can be found in [86].

The impact of the FALM has been studied in various works. [87] found that it improves the prediction of blade loads, especially near the root and tip, resulting in more accurate power output compared to the standard ALM. The wake structure remained largely unaffected.

[7] applied the FALM to a wing and wind turbine in laminar inflow. Their results confirmed that larger kernel sizes require stronger downwash corrections, especially near blade tips and roots. These corrections led to inward-shifted tip and root vortices, lower peak vorticity, and a slightly narrower shear layer. They also observed improved wake stability due to reduced velocity gradients in the tip vortex region.

4.3.2. Velocity Sampling

Accurate calculation of the blade load for body force projection requires the local relative velocity U_{rel} , which also determines the angle of attack used to retrieve aerodynamic force coefficients.

In blade element theory, this velocity is defined as the undisturbed inflow, meaning the velocity unaffected by the local aerodynamic forces at the blade section. While velocity data is readily available from the CFD simulation, identifying a sampling point that accurately represents the undisturbed flow remains a challenge.

In this work, the velocity is sampled at the cell centre nearest to each actuator point. However, due to the finite mesh resolution, this sampling location may not coincide with the centre of the bound vortex. As a result, the sampled velocity may include induction effects, introducing an error in the estimation of the freestream velocity.

4.3.3. Tower and Nacelle Modelling

Many ALM simulations, such as those by [33], model only blade forces. However, [88] demonstrated that inclusion of nacelle and tower effects is important for accurately capturing the near-wake aerodynamics. While a full geometric representation of these components in the CFD mesh is possible, it significantly increases computational cost and implementation complexity. Therefore, simplified body force models, similar to the actuator line method, are often used to represent their aerodynamic influence.

Tower

The tower can be modelled as an actuator drag line, discretising the tower into actuator points. At each point, the drag force per unit height is calculated using:

$$dF_D = \frac{1}{2} \rho U_\infty^2 C_d D_t dz \quad (4.20)$$

where ρ is the air density, U_∞ is the local freestream velocity at the tower, C_d is the drag coefficient, D_t is the tower diameter, and dz is the unit height. This force is then distributed over the surrounding mesh cells using a Gaussian kernel with a kernel size as used for the blades.

In AspFAST, an improved method applies Gaussian smoothing along the vertical (axial) and azimuthal directions, while assigning a constant force within the tower volume. This leads to the following kernel:

$$\eta(r, \theta, z) = \begin{cases} C e^{-\sqrt{\frac{z-z_0}{\epsilon}}}, & \text{if } r \leq R \\ C e^{-\sqrt{\frac{z-z_0}{\epsilon}}} e^{-\sqrt{\frac{r-R}{\epsilon}}}, & \text{if } r > R \end{cases} \quad (4.21)$$

where C is given by:

$$C = (R^2 \epsilon_z \pi^{1.5} + \epsilon^3 \pi^{1.5} + R \epsilon^2 \pi^2)^{-1} \quad (4.22)$$

As with blade force modelling, the velocity sampling location remains a challenge. [88] recommended sampling between the rotor and tower to capture rotor–tower interaction effects. However, sampling too close to the tower risks including induced velocity effects, leading to an inaccurate U_∞ .

In this study, the frontal area of a tower element is taken as the product of its height and diameter. The height depends on the number of actuator points and the diameter is equal to 8 cm. The drag coefficient is assumed to be 0.5, appropriate for a vertical cylinder in cross-flow [81].

This approach accounts only for drag forces acting on the flow. In reality, however, the tower also deflects the flow due to its impermeability and generates lift-induced effects such as vortex shedding. These phenomena are not captured by the simple drag-line model. To address this, more advanced modelling approaches are proposed in [88].

Nacelle

In AspFAST, a similar method as for the tower is used. The nacelle is represented by one node instead of multiple. The nacelle is modelled as a disk, normal to the rotor shaft direction. At the node, the force is calculated and distributed with the same method as for the tower such that inside the nacelle, the loading is constant. Outside the nacelle the load per cell reduces exponentially to zero.

The nacelle drag is characterised by its frontal area and drag coefficient. With dimensions of approximately 5.2 cm by 4.9 cm, the frontal area is about 26 cm² (leading to an equivalent radius). A drag coefficient of 1.05 is used, consistent with a cylinder aligned with the flow in low-speed regimes [81]. The velocity used for force calculation is sampled at a specified distance upstream of the nacelle actuator point to avoid local flow disturbances.

4.4. Hardware

The DelftBlue supercomputing cluster [89] is used to execute the simulations in this work. AspFAST requires one CPU and one GPU for each simulation, as OpenFAST runs on the CPU while GRASP is primarily GPU-based.

Due to memory constraints, the largest feasible simulation uses a mesh of approximately 225 million cells. Simulating 10 seconds of physical time (after the initial transient) takes roughly three days of wall-clock time.

4.5. Numerical-Experimental Comparison

This section describes how the results of the simulations will be compared with that of the experiments. First, the selected metrics will be discussed in Section 4.5.1, followed by where and at what rate the numerical results are acquired in Section 4.5.2.

4.5.1. Metrics

The numerical results will be assessed against experimental data, with a focus on key wake statistics. As discussed in Section 2.4.5, for sufficiently high Reynolds numbers, time-averaged quantities such as mean velocity, Turbulence Intensity (TI), vorticity, and Reynolds stresses become Reynolds number independent. This enables meaningful comparisons between scaled experiments and full-scale simulations of wind turbine wakes.

The time-averaged streamwise velocity is the primary metric, as it represents the bulk of kinetic energy in the wake. This is directly related to the power recovery potential for downstream turbines in wind farms. When speaking about the streamwise velocity, the terms recovery and velocity deficit occasionally enter the discussion. The terms refer to the recovery of kinetic energy as it was present in the inflow of the wind turbine. Velocity deficit is the velocity fraction that is 'missing' in the wake with respect to the inflow velocity.

The out-of-plane vorticity highlights wake instabilities that drive the transition from near wake to far wake, as outlined in Figure 2.2.1. For FOWTs, analysing vorticity offers insight into how these instabilities are affected by platform motion. Additionally, vorticity reveals regions where induced velocities significantly alter the streamwise velocity field.

The TI quantifies velocity fluctuations in the wake. These fluctuations are a key contributor to unsteady loads on downwind turbines, influencing their structural fatigue and lifetime. The definition of TI used here is directly related to the TKE, which is readily available in the experimental dataset:

$$TI = \left\langle \frac{U'_{rms}}{U} \right\rangle \quad (4.23)$$

with:

$$U'_{rms} = \sqrt{\frac{2}{3}k} = \sqrt{\frac{1}{3}(u'^2 + v'^2 + w'^2)} \quad (4.24)$$

$$U = \sqrt{\langle u \rangle^2 + \langle v \rangle^2 + \langle w \rangle^2} \quad (4.25)$$

Lastly, the Reynolds stresses (or kinematic shear stresses) describe the transport of momentum due to turbulent velocity fluctuations. For instance, considering a plane perpendicular to the y-axis (aligned with the turbine axis in Figure 1.1), the streamwise-vertical component $\langle u'w' \rangle$ represents transport of streamwise momentum in the vertical direction. The sign of this term is particularly important as it indicates the direction of energy transport within the wake.

All comparison metrics will be normalised for generality and to allow for direct comparison with experimental results. Velocities are normalised by the freestream velocity $U_\infty = 4 \text{ m/s}$, and lengths by the rotor diameter D . This applies to all velocity components, derived quantities such as vorticity and turbulence intensity, and spatial dimensions in plots and calculations throughout this work.

4.5.2. Output Acquisition

The flow field will not be evaluated across the full volumetric PTV dataset described in Chapter 3. Instead, a two-dimensional plane in the xz-direction, aligned with the rotor axis and the vertical direction, is used for comparison. This plane spans the domain $0.6 < x/D < 5.1$, $y/D = 0$ and $-0.12 < z/D < 0.65$, representing a cross-section through the measurement volumes shown in Figure 3.1.

In the planes, the metrics are evaluated along the vertical direction at eight streamwise positions: ($x/D = 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4$, and 5.1). Complete planes are included as contour figures in Appendix A. However, extracting profiles at specific locations enables more detailed comparisons between the different motion cases.

The experimental and numerical domains are both discretised. However, their discretisation is different. Therefore, a mismatch exists between the coordinates where data is provided. Inside the plane this challenge is overcome by linearly interpolating the experimental data to the coordinates of the numerical data. In the lateral direction, the mismatch is neglected, as the mesh resolution (chosen in Section 5.5) results in a maximal coordinate mismatch of 3.7% relative to the rotor diameter.

The sampling time of the wake is approximately 10 seconds, just like the experiments. An additional 5 seconds is added to let the wake transient pass. In this time the flow in the wake travels approximately two times the measured wake length. It has been verified in the results that starting after this timespan, the wake is not affected by the transient. Moreover, a sampling frequency of 50 Hz is maintained because this has been observed to be a good trade-off between a converging average and amount of data from one simulation.

Part III

Simulation Tuning

Parameter Study

This chapter addresses research question 1: “Which configuration options in AspFAST affect the wake development, and how?” Answering this is essential to establish a reliable base configuration, enabling subsequent investigations.

To establish this base configuration, a parameter study is performed, examining the key simulation parameters. The study is conducted under low-turbulence inflow conditions (TI below 0.1%) at rated wind speed. All simulations employ the ALM (not the FALM) and initially exclude tower and nacelle modelling.

The investigation begins by analysing inflow and outflow boundary conditions (Section 5.1), followed by an evaluation of domain size in the streamwise, lateral, and vertical directions, as well as the turbine’s placement along the streamwise axis (Section 5.2). The turbine is positioned centrally in the lateral and vertical directions to minimize boundary effects. Next, Section 5.3 explores the required number of actuator points on the blades. Subsequently, the impact of kernel size and grid resolution is assessed in Section 5.5. The final base configuration is then summarised in Section 5.6. The chapter concludes with a sensitivity analysis of turbulence generation in Section 5.8.

5.1. Inflow and Outflow Boundary condition

As discussed in Section 4.1.5, the inflow and outflow boundaries contain nudging zones to relax the cursor flow towards the precursor field, thereby maintaining a uniform inflow. This section evaluates the number of cells required in the nudging region at both the inflow and outflow to ensure sufficiently uniform velocity near the turbine inlet, while preserving as much physical domain as possible. To ensure a stable relaxation, the nudging factor is set to unity.

Figure 5.1 shows the normalised, time-averaged streamwise velocity just downstream of the inflow nudging zone along the vertical direction. The rotor hub is located at the origin of the domain, so $z/D = 0$ corresponds to the vertical position of the rotor axis and, assuming axisymmetry, the wake centreline. The profile is aligned with the lateral position of the rotor axis. Velocities are normalised by the freestream value of 4 m/s throughout this work. All simulations use identical settings except for the number of nudging cells at the streamwise boundaries.

The extent of the nudging region N_x is varied from 6 to 20 cells at the inflow and outflow in the streamwise direction. The figure shows that shorter nudging regions produce a noticeable velocity deficit. With 6 cells, the velocity is 98.75% of the freestream velocity, while 20 cells yield 99.9%. The velocity minimum follows an asymptotically decreasing trend with increasing nudging extent. Since an extent of 6 cells already results in a deficit of only 1.25%, this is deemed sufficient for the streamwise boundaries in this study.

5.2. Domain Size

The domain size depends on several factors. First, Section 5.2.1 examines the required offset between the turbine and the inflow boundary to allow the flow to adjust freely to the rotor forcing. Next, Section 5.2.2 investigates the distance from the rotor to the downwind boundary, as this boundary also effectively imposes a prescribed velocity and therefore pulls the wake. Finally, Section 5.2.3 addresses the required lateral and vertical dimensions to minimise blockage effects on the wake.




Figure 5.1: Normalised, time-averaged streamwise velocity at the end of the upwind nudging zone. The streamwise nudging extent is varied from 6 to 20 cells. The nudging factor is equal to 1 in all simulations.

5.2.1. Upwind Size

Flow approaching a wind turbine must decelerate gradually to satisfy the induction imposed by the rotor. This upwind induction zone should be influenced only by the turbine. However, the precursor method used at the inflow effectively acts as a Dirichlet boundary condition, prescribing a fixed velocity. If the turbine is placed too close to the inflow boundary, the flow cannot decelerate freely and results in an over-prediction of the velocity at the rotor. Therefore, the induction zone has been investigated to determine the required distance beyond the inflow nudging region.

Figure 5.2 shows the normalised, time-averaged, rotor-averaged streamwise velocity in the region upwind of the rotor, excluding the nudging zone. Here, rotor-averaged refers to a spatial average taken across the rotor diameter. The indicated distances exclude the upwind nudging zone.

The figure demonstrates that upwind distances of 0.3 D and 1.3 D deviate noticeably from the longer distances, both in the velocity at the rotor and in the gradient of the velocity curve. In contrast, the results for 1.9 D, 2.8 D, and 6.1 D are nearly identical. Therefore, an upwind distance of 2 D is considered sufficient for the base configuration.

5.2.2. Downwind Size

In the experiments the wake was measured up to 5.1 D after the rotor plane. Therefore, the simulations must include the flow at least until this distance. Additionally, extra space is needed beyond this to prevent boundary effects, as the prescribed outflow velocity can artificially accelerate the wake. This non-physical acceleration makes a portion of the downstream domain unusable for analysis.

To evaluate this effect, four simulations with varying downwind extents were conducted. Similar to Figure 5.2, Figure 5.3 presents the normalised, time-averaged, rotor-averaged streamwise velocity downstream of the rotor, excluding the outflow nudging zone. The downwind distances indicated in the legend exclude the outflow nudging zone as well.

At the rotor plane ($x/D = 0$), the velocity is 75% of the freestream value. After the wake expansion a velocity minimum of 50% of the freestream is reached. This behaviour is consistent with one-dimensional actuator disc theory, which predicts that the velocity in the far wake is approximately half that at the rotor [3].

Beyond this point, the simulation with a 10.7 D downwind extent begins to show a steep velocity increase due to outflow effects. The other cases also show signs of wake recovery but at a more gradual rate. Minor differences among these curves are attributed to slight variations in the inflow; however, their overall trends remain consistent.

At $x/D = 15$, the 16 D downwind distance simulation start to diverge from the 20 D and 25.3 D downwind distance simulations. Subsequently, the 20 D starts to differ from the simulation with the longest downwind distance at $x/D = 19$.

These results show that the final 2 D of the domain is typically affected by the outflow boundary. This influence diminishes when the wake has sufficiently recovered, as less artificial acceleration is needed to match the prescribed outflow velocity.

In this work, an additional 5 D is added beyond the experimental region of interest, resulting in a total downwind extent of 10 D. This choice is conservative, as earlier results suggested a larger required distance. However, that conclusion was influenced by the time-averaging interval including part of the wake transient at the start of the simulation.




Figure 5.2: Normalised, time-averaged, rotor-averaged streamwise velocity, upwind of the rotor. The five curves exclude the upwind nudging zone and show that the upwind velocity converge to one velocity development as the upwind distance increases.




Figure 5.3: Normalised, time-averaged, rotor-averaged streamwise velocity, downwind of the rotor. The four curves exclude the downwind nudging zone and show that approximately 2 D of downwind domain is unphysical due to the influence of the outflow boundary.

5.2.3. Lateral and Vertical Size

As explained in Equation 4.1.5, lateral flow in the simulation imposes pressure on the interior of the domain, resulting in the so-called blockage effect. This effect leads to reduced wake expansion and an overprediction of velocity due to a smaller effective cross-sectional area. Preventing blockage is therefore essential to accurately capture wake characteristics, much like in wind tunnel testing.

One way to quantify the blockage effect is through the turbine's performance, as an overestimated velocity leads to a higher thrust coefficient. The thrust coefficient is thus analysed in relation to the Blockage Ratio (BR), defined in Equation 5.1:

$$BR = \frac{A_{rotor}}{A_{tot}} \quad (5.1)$$

Figure 5.4 shows the thrust coefficient plotted against the blockage ratio. The blue dots represent simulated values, while the green dashed line is a linear fit through these points. A clear linear decrease in thrust coefficient is observed as the blockage ratio decreases. Extrapolating the trend to zero blockage yields a thrust coefficient of 0.8318, which aligns well with experimental results [51].

However, zero blockage would require an infinitely large domain, which is computationally unfeasible. Moreover, it is not necessary to eliminate blockage entirely. As discussed in Equation 4.1.5, keeping the blockage ratio below 5% is sufficient to minimise its influence on the wake. This corresponds to a cross-sectional domain size of at least $4.5 D \times 4.5 D$ in width and height.

The effect of blockage is also evident in the velocity profile shown in Figure 5.5. This figure presents the normalised, time-averaged streamwise velocity along the vertical direction at $x/D = 2$. The profile is

aligned with the lateral position of the rotor axis. It can be observed that higher blockage ratios lead to an overall speed-up, both within the wake core and near the edges of the wake.

When the blockage ratio is reduced below 4.4%, corresponding to a domain cross-section of $4.8 D \times 4.8 D$, the velocity profile shows little change compared to cases with even lower blockage. Based on this, a cross-section of $5 D \times 5 D$ is deemed sufficient for the present study.

In addition to the inflow and outflow boundaries, the lateral boundaries are treated as periodic. Since no wake enters or exits through these sides and the turbine is located well away from them, only three nudging cells are applied in these directions.

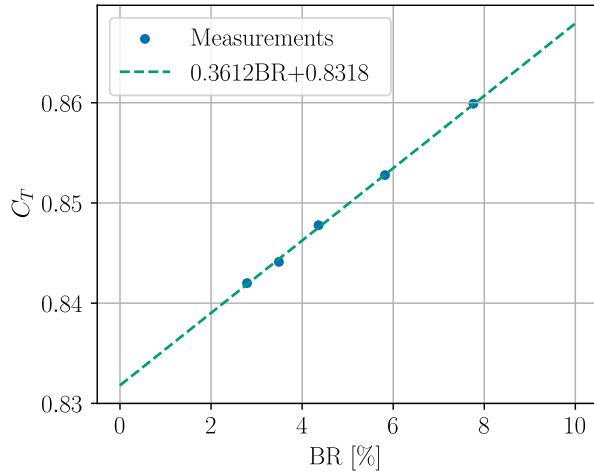


Figure 5.4: Thrust coefficient C_T against the blockage ratio, indicating a linear reduction of thrust as blockage reduces.




Figure 5.5: Normalised, time-averaged velocity along the vertical direction at $x/D = 2$. The velocity converges for low blockage ratios.

5.3. Actuator Points

As discussed in Chapter 4, literature suggests using between 30 and 60 actuator points along the blade span, while also ensuring this number does not exceed the number of grid cells along the blade [83]. Increasing the number of actuator points incurs a computational cost, as each point requires serial processing. Therefore, the number of points should be kept to the minimum necessary for accurate results.

Figure 5.6 presents the normalised, time-averaged streamwise velocity profile at $x/D = 2$ for four simulations with increasing numbers of actuator points. The profile is aligned with the lateral position of the rotor axis. The case with 30 actuator points shows a clear deviation near $z/D = 0.25$, and the case with 40 points shows a small difference at the velocity minimum near $z/D = 0.375$. However, the curves for 50 and 60 points are nearly identical, indicating that at least 50 points are needed for the solution to become independent of this parameter.

5.4. Kernel Size

As explained in Section 4.3, the kernel size should at least be 2 times the cell size to prevent spurious oscillation in the flow. However, this is not necessarily enough. Therefore, the wake should be investigated for a varying kernel size.

In Figure 5.7 the normalised, time-averaged streamwise velocity along the vertical direction at $x/D = 2$ is presented for values of $\epsilon/\Delta x$ equal to 2, 3 and 5. The lowest kernel size results in the highest velocity gradients in the root and tip regions of the wake, while a higher size causes the wake to diffuse. For $\epsilon/\Delta x = 2$ and low resolution, there is a small fluctuation around $z/D = 0.4$. This is solved by increasing the kernel size to $\epsilon/\Delta x = 3$. Additionally, due to the relation between the kernel size and the grid resolution, a $\epsilon/\Delta x$ of 2 becomes even worse for a high grid resolution. This can be seen in the curve for a resolution $\Delta x/D = 0.012$. In this curve, many fluctuations can be observed due to a too low absolute kernel size.

Since $\epsilon/\Delta x = 3$ does not contain spurious fluctuation in the wake, while limiting the decrease of the velocity gradients in the root and tip regions of the blades, this is the most optimal kernel size and will be used in this work.

As explained in Section 4.3, the kernel size ϵ must be at least twice the local grid spacing Δx to prevent spurious oscillations in the flow. However, this minimum is not necessarily sufficient for accurate results, so the wake behaviour should be evaluated for varying kernel sizes.

Figure 5.7 shows the normalised, time-averaged streamwise velocity along the vertical direction at $x/D = 2$, for values of $\epsilon/\Delta x$ equal to 2, 3, and 5. The lowest kernel size $\epsilon/\Delta x = 2$ results in the steepest velocity gradients in the root and tip regions of the wake, while larger values of ϵ cause increased diffusion. For $\epsilon/\Delta x = 2$, a small fluctuation appears around $z/D = 0.4$, which is resolved when increasing the kernel size to 3.

Furthermore, because the kernel size is defined relative to the grid spacing, a value of $\epsilon/\Delta x = 2$ becomes more problematic at higher resolutions, where the absolute kernel size becomes too small. This is evident in the curve for a grid resolution of $\Delta x/D = 0.012$, where spurious fluctuations are visible throughout the wake. Since $\epsilon/\Delta x = 3$ removes these spurious fluctuations while preserving the velocity gradients near the blade root and tip, it is deemed optimal and is adopted for all simulations in this work.




Figure 5.6: Normalised, time-average streamwise velocity profile along the vertical direction at $x/D = 2$ for different numbers of actuator points. The profile converges for 50 actuator points on each blade.




Figure 5.7: Effect of reducing the kernel size ϵ on the normalised vertical velocity profile at $x/D = 2$.

5.5. Grid Resolution

According to [25], a well-functioning LES with wall modelling should resolve at least 80% of the total TKE. Although, the simulation performs no wall modelling, several resolutions have been tested for this criterion.

Figure 5.8 shows the spatial distribution of the time-averaged resolved TKE fraction for three mesh resolutions: $\Delta x/D = 0.017$, 0.012 and 0.011. Although the fields display some fluctuations, caused by undersampling and a relatively short averaging interval, they show the expected trend: regions containing the root and tip vortices show a low resolved TKE fraction. This is the result of the high turbulence production due to a high velocity gradient. However, the produced turbulence scales are relatively small and are therefore not resolved by the mesh. As a result, none of the cases fully meet Pope's 80% resolution criterion.

These and additional coarser simulations were used to develop a fitted function for the time-averaged, domain-averaged ratio of resolved to total TKE, as shown in Figure 5.9. Here, the fraction is averaged spatially from the rotor plane towards the effective outflow (excluding the last 5 D as discussed before), between $z/D = \pm 1$. From this, it can be estimated that a resolution of $\Delta x/D = 0.005$ would be required

to satisfy the criterion. However, such a resolution would require approximately 258 million cells for the current domain, which is computationally unfeasible. For more on computational resource limitations, refer to Section 4.4.

The primary reason for not meeting Pope’s criterion is the absence of turbulence at the inflow. In the current setup, large, energetic turbulent structures are lacking, which means that a substantial portion of the total TKE is concentrated in the small-scale structures generated within the shear layers of the wake. These structures are not resolved by the mesh, leading to a low resolved TKE fraction.

In Chapter 8, turbulence will be introduced at the inflow. The inflow turbulence structures are both larger and more energetic, making them more likely to be resolved by the mesh. While the small-scale turbulence in the root and tip shear layers will still be present, their relative contribution to the total TKE will diminish. As a result, the resolved fraction of TKE will increase due to the greater weight of the larger, resolved inflow structures.




Figure 5.8: Fraction of resolved TKE with respect to the total TKE in the near wake for varying mesh resolution. The resolutions $\Delta x/D$ from top to bottom are 0.017, 0.012 and 0.011

Given the limitation of Pope’s criterion in low-turbulence inflow, the normalised, time-averaged stream-wise velocity along the vertical direction at $x/D = 2$ is presented in Figure 5.10 for various resolutions ranging from $\Delta x/D$ between 0.041 and 0.012.

The results show increasing velocity gradients in the tip and root regions with higher resolution. The minimum velocity remains consistently around $z/D = 0.3$, though its local curvature flattens with increasing resolution. From $\Delta x/D = 0.017$ onwards, there is even a small oscillation present between $z/D = 0.1$ and 0.4, which indicates an insufficient kernel size. However, this will not be changed since increasing it will result in more diffused shear layers.

The region containing the oscillations corresponds to the ϵ/c minimum presented in Figure 5.12. Local mesh refinement at the root and tip regions would allow for high velocity gradients while preventing oscillations in the wake core.

Around the rotor axis ($z/D = 0$), higher velocities are observed, attributed to reduced force smearing in the region beyond the physical blade root. Also, the case with the highest resolution presents a small velocity minimum at the wake axis. This is the effect of nacelle drag and should be ignored at this point of the analysis.

The curves for $\Delta x/D = 0.017$ and 0.012 are very similar, indicating that the solution is approaching convergence. Since the finer case requires approximately twice the computational resources, $\Delta x/D = 0.017$ is selected as the base resolution for further simulations, offering a balance between accuracy and efficiency for parameter tuning.

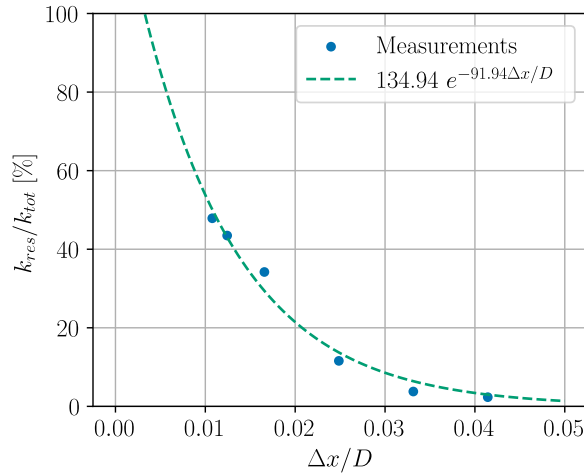


Figure 5.9: Time-averaged, space-averaged ratio of resolved and total TKE averaged from the turbine rotor plane to the domain outflow and between $z/D = \pm 1$ along the vertical direction. The trend through the measurements shows an exponential fit.




Figure 5.10: Normalised, time-averaged streamwise velocity along the vertical direction at $x/D = 2$ for varying resolutions. A higher resolution increases the velocity gradients, but could cause an insufficient kernel size.

5.6. Base Case

The parameter study leads to a well-defined base case configuration. To ensure a maximum velocity deficit of 1.25% in the nudging zones, both the upwind and downwind zones should be six cells wide. Additionally, a minimum distance of 2 D between the inflow boundary and the rotor plane is required to allow for natural deceleration of the flow prior to reaching the turbine. No turbulence is synthesised in the inflow.

To obtain a wake length of 5 D for analysis while limiting the influence of the outflow nudging zone, a total downstream distance of 10 D from the rotor to the outflow boundary is necessary. In the lateral and vertical directions, the domain should extend at least 5 D in width and height, respectively, with the turbine hub placed at the domain centre to minimise blockage effects.

For wake convergence, a grid resolution of $\Delta x/D = 0.017$ is required. Due to GPU-based computational constraints, the final domain dimensions are adjusted to 12.2 D in length, and 5.3 D in both width and height.

Figure 5.11 illustrates the resulting computational domain. The gray regions indicate the nudging zones, and the red region marks the area lost due to the influence of the outflow boundary. The turbine is also shown, with a black bullet representing the centre of rotation for pitch motion (blue arrow) and the neutral position for surge motion (red arrow).

The base simulation employs the conventional ALM. A kernel size of $\epsilon/\Delta x = 3$ is used to suppress spurious fluctuations in high-resolution meshes while maintaining adequate shear layer representation. A total of 50 actuator points is selected, which has been shown to be sufficient for solution convergence. The nacelle and tower are excluded from the base case. Finally, the second-order central difference advection scheme is utilised, and the Smagorinsky turbulence model is applied, using a model constant of 0.17, following the recommendation in [25].

5.7. FALM

As discussed in Section 4.3, the optimal kernel size ϵ is ideally one quarter of the local chord length. However, Figure 5.12 shows that the ratio ϵ/c deviates significantly from this ideal value at many spanwise positions, even for a relatively high resolution (green curve).

In GRASP, the kernel size is defined as a fixed ratio $\epsilon/\Delta x$, and a uniform grid is used. This makes it impossible to match the optimal ϵ/c across the full blade span. As a result, the ALM cannot accurately reproduce the local aerodynamic response, particularly in regions with coarse resolution.




Figure 5.11: Schematic of the domain. The white area is physically useful, while the gray area is affected by the boundary conditions. Left is a front view and right is a side view. The turbine point of rotation is indicated by the black bullet. The red arrow indicates the surge motion and the blue arrow indicates the pitch motion.

To mitigate this limitation, the FALM introduces a correction to the downwash induced by the actuator forces, compensating for the smoothing effect of a large kernel size. The impact of applying the FALM is illustrated in Figure 5.13, which shows the normalised, time-averaged streamwise velocity profiles at $x/D = 2$ for two kernel sizes ($\epsilon/\Delta x = 2$ and 5) at low resolution. Solid lines represent results obtained with the standard ALM, while dashed lines show the same configurations using the FALM.

The application of the FALM results in higher velocity throughout the wake, especially in the tip and root regions. This reflects the enhanced downwash, which compensates for the under-resolved tip vortex in coarse meshes. Notably, this effect is present regardless of kernel size.

As shown in [7], the influence of the FALM diminishes as resolution increases, since a smaller ϵ/c (for a finer mesh) reduces the need for correction. Therefore, while the FALM is critical at low resolutions, its impact becomes negligible for finer grids.

5.8. Turbulence Synthesis

Although not included in the base configuration, turbulence in the inflow is important for accurate wake development [59]. In GRASP, turbulence is generated through a flux of humidity and/or heat. On the other hand, in the wind tunnel, this mechanism is not physically present, so it is not essential which of the two is used. Therefore, heat flux has been selected as the turbulence generation mechanism in this work.

Turbulence synthesis (excluding the shear stresses) is governed by two parameters: the perturbation magnitude A_θ , and the heat flux $w\theta_T$. The former sets an initial disturbance in the flow field, while the latter introduces random temperature fluctuations through the bottom boundary over time.

The definition of the TI is provided in Section 4.5 and is analysed in Figure 5.14a and Figure 5.14b. Spatial averaging is performed from the inflow to the effective outflow and radially across the rotor span, excluding boundary regions. Temporal averaging is done over 10-second intervals to match the measurement duration in the experiment.

5.8.1. Heat Perturbation Magnitude

Figure 5.14a presents the turbulence intensity for varying perturbation magnitudes with a constant heat flux. The initial value of TI increases approximately linearly with A_θ , as shown in Figure 5.14c. Following the initial peak, all curves drop rapidly to a minimum around 1.25%, with the timing of this minimum being delayed at higher perturbation magnitudes.

Subsequently, TI increases again, with a similar rate across all cases. However, both the timing and height of the resulting maximum vary: higher perturbation magnitudes delay the peak and result in lower maximum values. After reaching this maximum, the TI oscillates in time.




Figure 5.12: Kernel size with respect to the local chord length for two kernel sizes. Both have an $\epsilon/\Delta x$ of 2, but with resolutions $\Delta x/D$ of 0.041 and 0.012.




Figure 5.13: Effect of the FALM on the normalised, time-averaged streamwise velocity. The blue curves have a relative kernel size of 2 and the green curves one of 5, with equal resolutions.

5.8.2. Heat Flux

Figure 5.14b shows the influence of the heat flux. All curves begin with similar TI values due to a shared perturbation magnitude. Approximately 25 seconds are required before an increase in TI is observed. This delay reflects the time needed for turbulence generated at the surface to reach the rotor height.

The growth rate of TI is clearly influenced by the magnitude of the flux: larger heat fluxes result in steeper increases. Each curve reaches a peak, the height of which increases with increasing flux. After the peak, the turbulence intensity enters a fluctuating regime.




Figure 5.14: Effect of heat perturbation magnitude and heat flux on the TI.

Experimental Results

This chapter analyses the experimental results to establish the baseline wake characteristics across all motion cases: stationary, Low-Frequency Surge (LFS), High-Frequency Surge (HFS), Low-Frequency Pitch (LFP), and High-Frequency Pitch (HFP). The terms “low-” and “high-frequency” correspond to the 2 Hz and 5 Hz motion cases, respectively, as described in Chapter 3. Furthermore, the influence of each motion case is evaluated relative to the stationary turbine configuration. At the end of this chapter, research question 2 will be answered.

First, the experimental results are evaluated in terms of the time-averaged streamwise velocity, out-of-plane vorticity, turbulence intensity, and streamwise-vertical Reynolds stresses, as presented in Section 6.1, Section 6.2, Section 6.3, and Section 6.4, respectively. Additionally, the vertical velocity field is examined, as it offers relevant insight into the discrepancies between simulations and experimental observations in Chapter 8. Finally, the observations in the metrics are used to understand their dependencies on each other in Section 6.6.

6.1. Streamwise Velocity

In this section, observations in streamwise velocity of the wake will be discussed. Figure 6.1 presents the normalised time-averaged streamwise vorticity along the vertical direction at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. The wake has been divided into four main characteristics and regions: the root region, wake core, tip region and far wake. The profiles are dashed lines. This convention for the experimental results is used throughout this work. Solid lines will be used for numerical results.

6.1.1. Root Region

At the start of the measured domain, all motion cases display the same key feature: a nacelle jet between $z/D = -0.05$ and $z/D = 0.05$, caused by the nacelle’s low aerodynamic loading. The jet is asymmetric about $z/D = 0$ because the tower reduces streamwise velocity in the lower part; this effect is largely gone by $x/D = 1.3$. The streamwise velocity peaks at $z/D = 0.05$ with a normalised value of 0.8.

Moving downstream, the nacelle jet streamwise velocity decreases: first through viscous and turbulent diffusion while the root vortices remain intact, and later—after vortex breakdown—through enhanced turbulent mixing. By $x/D = 1.3$, the tower’s influence is barely visible, and the jet disappears entirely before $x/D = 3.1$, where its maximum streamwise velocity has dropped to 0.5.

Differences between motion cases in this region are modest but noticeable. The Low-Frequency Surge (LFS) case shows the fastest streamwise velocity reduction by $x/D = 2.5$. The High-Frequency Surge (HFS) and Low-Frequency Pitch (LFP) cases also decelerate quickly, while the stationary and High-Frequency Pitch (HFP) cases remain similar to each other and retain higher velocities further downstream.

The root vortex shear layer is a second prominent feature, marked by a sharp negative streamwise velocity gradient between $z/D = 0.05$ and $z/D = 0.20$. Initially this gradient is produced by the induced streamwise velocity of the root vortex; following breakdown the strong mean gradient generates turbulence, which diffuses the shear layer. At $x/D = 0.7$ the motion cases are similar, but by $x/D = 1.3$ all motion cases except HFS show faster shear-layer diffusion than the stationary case. By $x/D = 1.9$ and especially at $x/D = 2.5$, the streamwise velocity gradients for moving turbines are notably smaller than for the stationary turbine, suggesting faster diffusion when the turbine is in motion.

6.1.2. Wake Core

The third region, the wake core, corresponds to the flow generated by the mid-span of the blades and extends from $z/D = 0.2$ to $z/D = 0.40\text{--}0.45$, with a normalised streamwise velocity between 0.35 and 0.40. At $x/D = 0.7$, its width depends on the motion case: the pitch cases and the HFS case show a wake core that is approximately $0.05 D$ wider than in the stationary case, a difference attributable to the vorticity present in this region. In contrast, the surge cases have a wake core height similar to that of the stationary turbine. Pitch motion also shifts the streamwise velocity minimum toward the tip, and all motion cases exhibit a slight streamwise velocity increase relative to the stationary case in the lower portion of the wake core ($z/D < 0.3$).

Downstream, the wake core streamwise velocity decreases to about 0.25 by $x/D \approx 1.6$. Beyond this point, the core begins to contract and accelerate, marking the onset of wake recovery. Low-frequency motion cases recover fastest, followed by high-frequency cases; both recover more quickly than the stationary turbine. At $x/D = 2.5$, the LFP case shows a greater recovery than the LFS case, with a maximum difference of 0.05 in normalised streamwise velocity.

6.1.3. Tip Region

At the top of the wake, the tip vortex shear layer extends from the edge of the wake core to $z/D = 0.6$. This region is characterised by a strong positive streamwise velocity gradient induced by the tip vortex. Initially, the gradient is similar across all motion cases, corresponding to a normalised streamwise velocity range of 0.4 to 1.0.

After reaching the streamwise velocity minimum at $x/D \approx 1.6$, the gradient decreases rapidly, indicating the onset of tip vortex breakdown. Low-frequency motion cases exhibit the greatest additional diffusion relative to the stationary case, while high-frequency cases also enhance diffusion but to a lesser extent. At $x/D = 2.5$, a notable difference is observed between the low-frequency cases, consistent with trends in the wake core, whereas the high-frequency cases remain similar. The difference between the LFS and LFP cases disappears by the far-wake transition at $x/D = 3.1$.

6.1.4. Far Wake

The middle measurement volume ($x/D = 3.1\text{--}3.8$) marks the transition from the near wake to the far wake. By $x/D = 3.1$, the root shear layer has fully diffused. Progressing to $x/D = 3.8$, the wake continues to develop towards an axisymmetric form as the shear layer further diffuses.

Differences between motion cases show that low-frequency motions recover the fastest, with the LFS case overtaking the LFP case in recovery rate. Due to this accelerated recovery, both low-frequency cases already display a Gaussian-like profile at $x/D = 3.1$. High-frequency cases, along with the stationary case, maintain a nearly constant normalised streamwise velocity of 0.45 up to $z/D = 0.2$. By $x/D = 3.8$, the relative ordering of cases remains the same as at $x/D = 3.1$, but overall velocities have increased. At the wake axis, low-frequency motion still produces the highest recovery, followed by the high-frequency cases, both outperforming the stationary turbine.

After the second gap between measurement volumes, the LFS case exhibits an increase in streamwise velocity at the wake axis. For the LFP case, this change is minimal. In contrast, the high-frequency cases show a significant velocity increase between the wake axis and $z/D = 0.4$, producing an almost linear velocity rise with height. Among all motion cases, the HFP case has the lowest recovery above $z/D = 0.4$, even lower than the stationary case. Interestingly, the HFS case matches the LFP case in velocity from $x/D = 4.4$ onwards due to a pronounced speed-up in its wake.

6.2. Vorticity

Closely related to the velocity field is the vorticity generated by the blades as explained in Figure 2.2.1. Therefore, Figure 6.2 presents the normalised time-averaged out-of-plane vorticity along the vertical direction at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

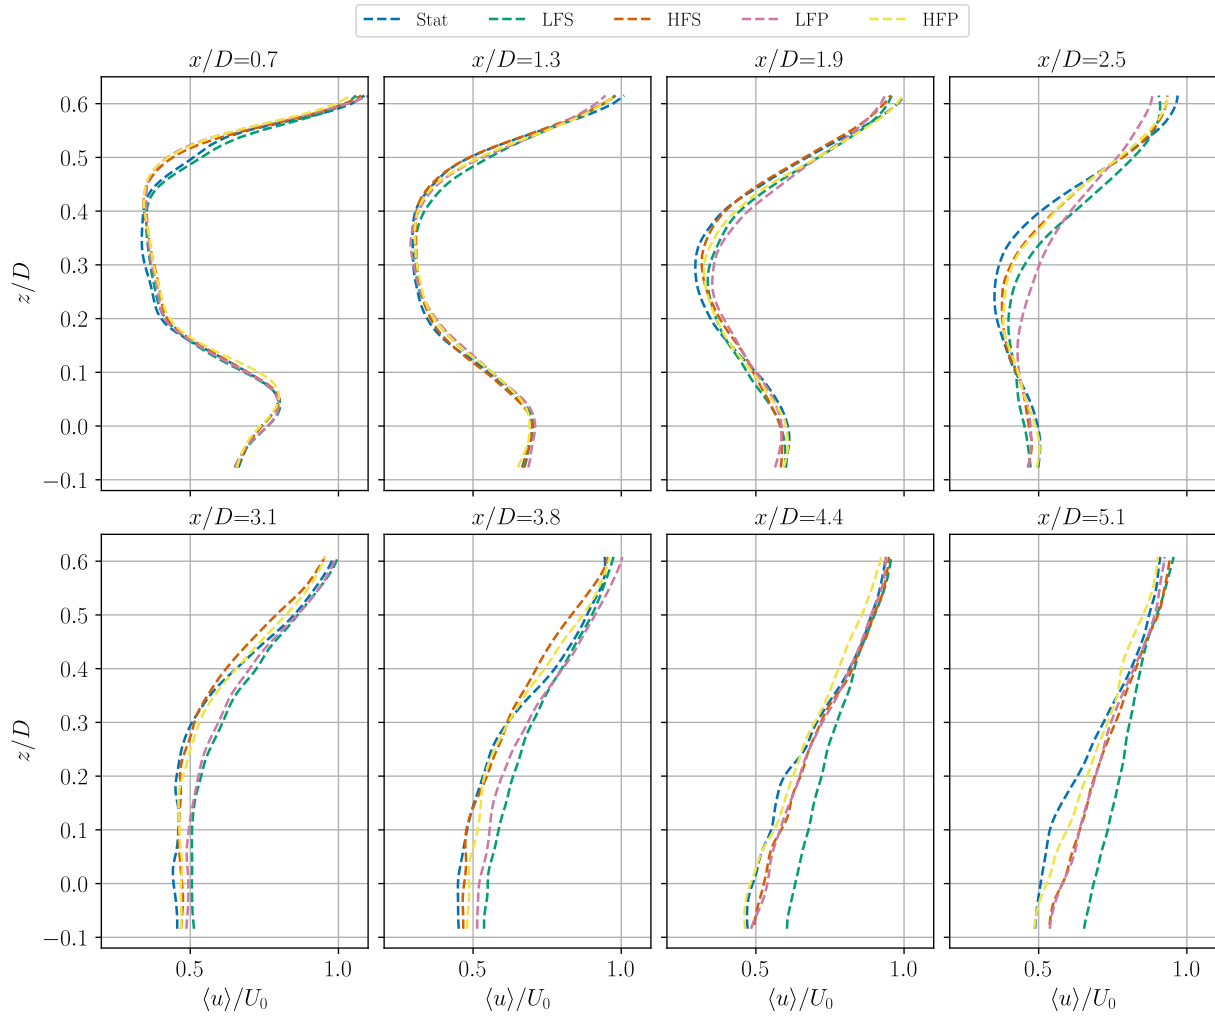


Figure 6.1: Normalised time-averaged streamwise velocity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

6.2.1. Root Region

At the start of the domain, the nacelle jet region, located between $z/D = -0.05$ and $z/D = 0.05$, contains a vorticity maximum at $z/D = 0$. This peak corresponds to vorticity generated by the strong force gradient below the hub. Notably, this gradient originates at the blade root section, which is positioned below the hub rather than at its height. A likely explanation is that the tower induces an upward displacement of the lower root vortex. This vorticity maximum diffuses rapidly, reaching zero out-of-plane vorticity by $x/D = 1.3$ for all motion cases. Such rapid decay could be linked to the destabilising influence of the tower and nacelle [23].

A vorticity minimum at $z/D = 0.1$ marks the upper part of the root vortex system. Its absolute magnitude exceeds that of the maximum at the wake axis, likely due to faster diffusion of the lower vortex half caused by tower and nacelle interactions. The root vortex diffusion lasts until $x/D = 2.5$, by which point it has nearly vanished.

Although differences between motion cases are small, some trends emerge. At $x/D = 0.7$, the LFS, HFS, and LFP cases exhibit lower vorticity than the stationary and HFP cases. By $x/D = 1.3$, the HFP case retains more concentrated root vorticity than the others, suggesting a more stable vortex system due to its motion. The remaining cases show no significant differences relative to each other.

6.2.2. Wake Core and Tip Region

Initially, the wake core contains no vorticity due to the irrotational inflow field. Further downstream, the tip vorticity diffusion cause the wake core to narrow, with the tip vortex producing a positive vorticity maximum. Notable differences exist between motion cases: at $x/D = 0.7$, the high frequency cases exhibit an increase in vorticity, while the low-frequency cases show a vorticity decrease, especially the LFS case.

Both the stationary and LFS cases display a sudden change in vorticity gradient between $z/D = 0.4$ and $z/D = 0.5$, attributed to secondary vortices at this vertical position [6]. These vortices are weak and do not significantly destabilise the flow.

Beyond $x/D = 1.6$, the tip vortex undergoes pronounced breakdown, causing the vorticity maximum to diffuse. Differences among motion cases diminish, although up to $x/D = 2.5$, low-frequency cases retain a more diffused tip vortex remnant. High-frequency cases are also more diffused than the stationary case. After $x/D = 3.8$, vorticity values converge near zero and fluctuate slightly, maintaining this uniform state until the end of the measurement domain.

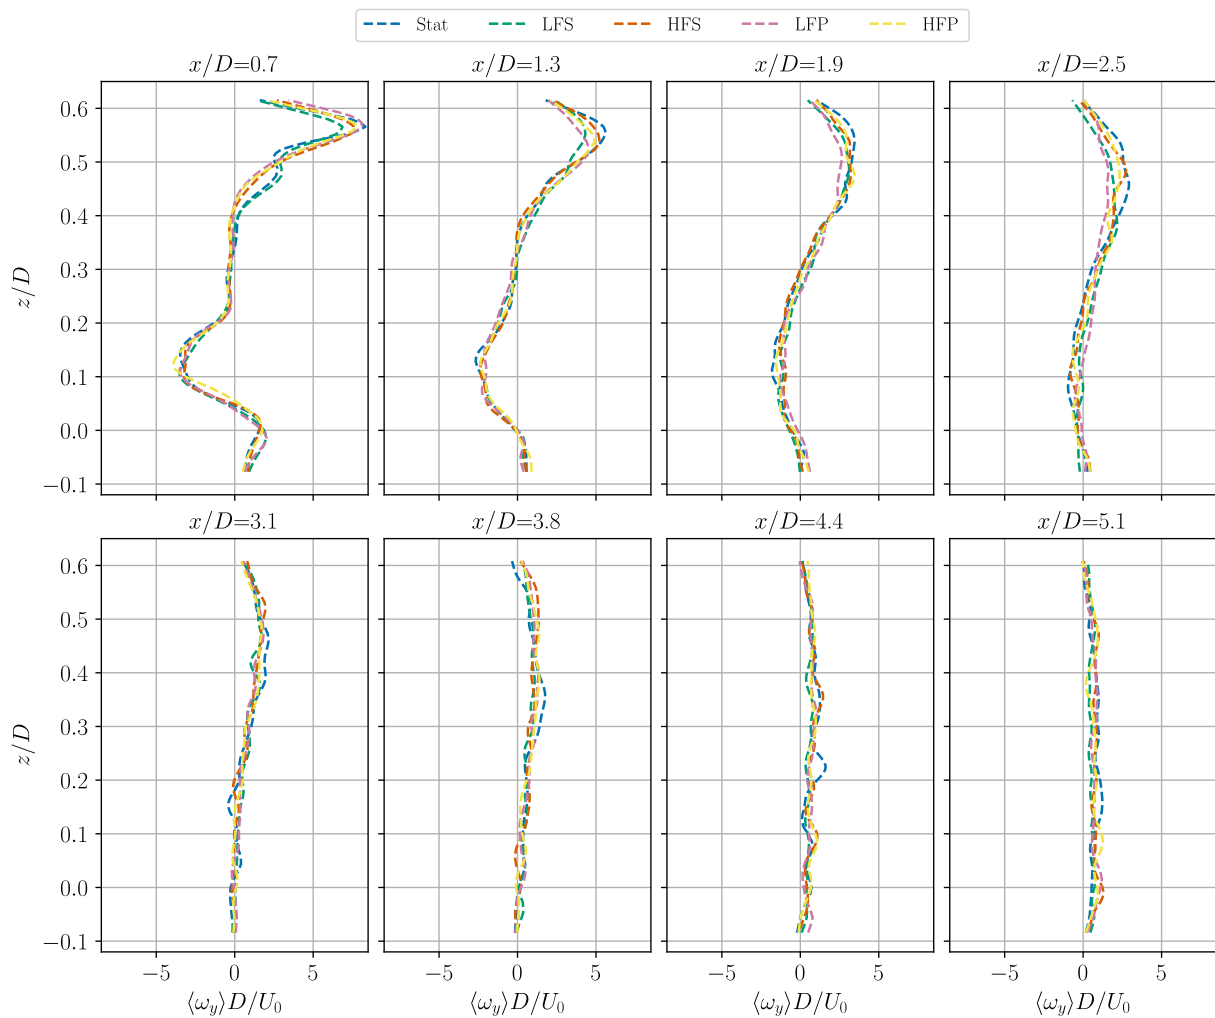


Figure 6.2: Normalised time-averaged out-of-plane vorticity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

6.3. Turbulence Intensity

Turbulence contributes to the entrainment of the wake. As was explained in Figure 2.2.1, much turbulence is produced when the helical tip vortex structures breakdown. It was shown that the floater motion has an effect on these structures. Therefore, studying the turbulence is important too. For this reason, Figure 6.3

presents the time-averaged TI along the vertical direction at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

6.3.1. Root Region

At $x/D = 0.7$, below $z/D = 0.2$, the TI exhibits two local maxima and one minimum. A maximum of 12% is observed at $z/D = 0.15$, just above the local minimum in out-of-plane vorticity associated with the upper part of the root vortex. A minimum of approximately 10% TI occurs at $z/D = 0.05$, corresponding to the zero-crossing of out-of-plane vorticity. The second maximum, around 11%, is located at $z/D = -0.05$, just below the out-of-plane vorticity maximum at $z/D = 0$.

Among the cases, the high-frequency motions exhibit the strongest fluctuations, with $\sim 1\%$ higher TI compared to the stationary case. Low-frequency cases follow with about 0.5% additional turbulence intensity. The stationary turbine shows the lowest TI in the root vortex region at $x/D = 0.7$.

By $x/D = 1.3$, the lower maximum disappears, replaced by a minimum at the wake axis, indicating an emerging axisymmetric trend likely related to the diminishing influence of the tower and nacelle. The upper maximum at $z/D = 0.05$ persists with little change in TI. The low-frequency cases exhibit a more diffused maximum, suggesting significant root vortex breakdown, whereas the stationary and high-frequency cases maintain a distinct maximum.

6.3.2. Wake Core and Tip region

At $x/D = 0.7$, the wake core exhibits the absolute minimum turbulence intensity with notable differences between motion cases. The stationary case shows a minimum TI of 5%, while the low-frequency cases reach 7%, and the high-frequency cases 8%. This reflects that faster turbine motion induces stronger fluctuations in the flow near the turbine, unlike the stationary turbine.

By $x/D = 1.3$, the differences between motion cases increase: the stationary case maintains the lowest TI at 6%, the HFS case narrows the gap to 8%, the LFS and HFP cases both show TI values around 10%, and the LFP case reaches the highest TI of 12%.

At $x/D = 1.9$, the low-frequency cases no longer display a low-turbulence wake core due to breakdown of the root and tip vortices. Conversely, the stationary and high-frequency cases still retain signatures of the turbulence maximum in the root region, resulting in a small minimum in turbulence intensity within the wake core region. Consequently, their TI levels have not increased as much as those of the low-frequency cases at this downstream position.

6.3.3. Far Wake

The TI increases notably toward $x/D = 2.5$, reaching a maximum of 34% for the LFS case. Up to this point, the wake evolution is similar across all motion cases, although a delay is observed in the stationary and high-frequency cases compared to the low-frequency ones.

Further downstream, the TI in the low-frequency cases begins to decrease, dropping to about 28% above the wake axis at $x/D = 3.1$. This marks the onset of a Gaussian profile, indicative of an axisymmetric wake developing beyond $x/D = 3.1$. Within these profiles, the LFS case starts to lag behind other motion cases, especially near the wake edge above $z/D = 0.4$. This trend arises partly from a reduction in LFS turbulence intensity, while other cases experience a TI increase up to $x/D = 4.4$.

Subsequently, the LFP case also falls behind the LFS in terms of TI evolution, resulting in a 5% difference around the wake axis, though no difference is noted above $z/D = 0.1$. The high-frequency cases display similar behaviour, with about a 1% increase in TI near the wake axis. Additionally, the HFP case shows higher TI at the top of the wake compared to other cases. The stationary turbine exhibits the largest delay in TI development. Overall, these observations suggest that turbine motion initially promotes turbulence generation in the wake, but at farther downstream locations, motion leads to an earlier turbulence reduction.

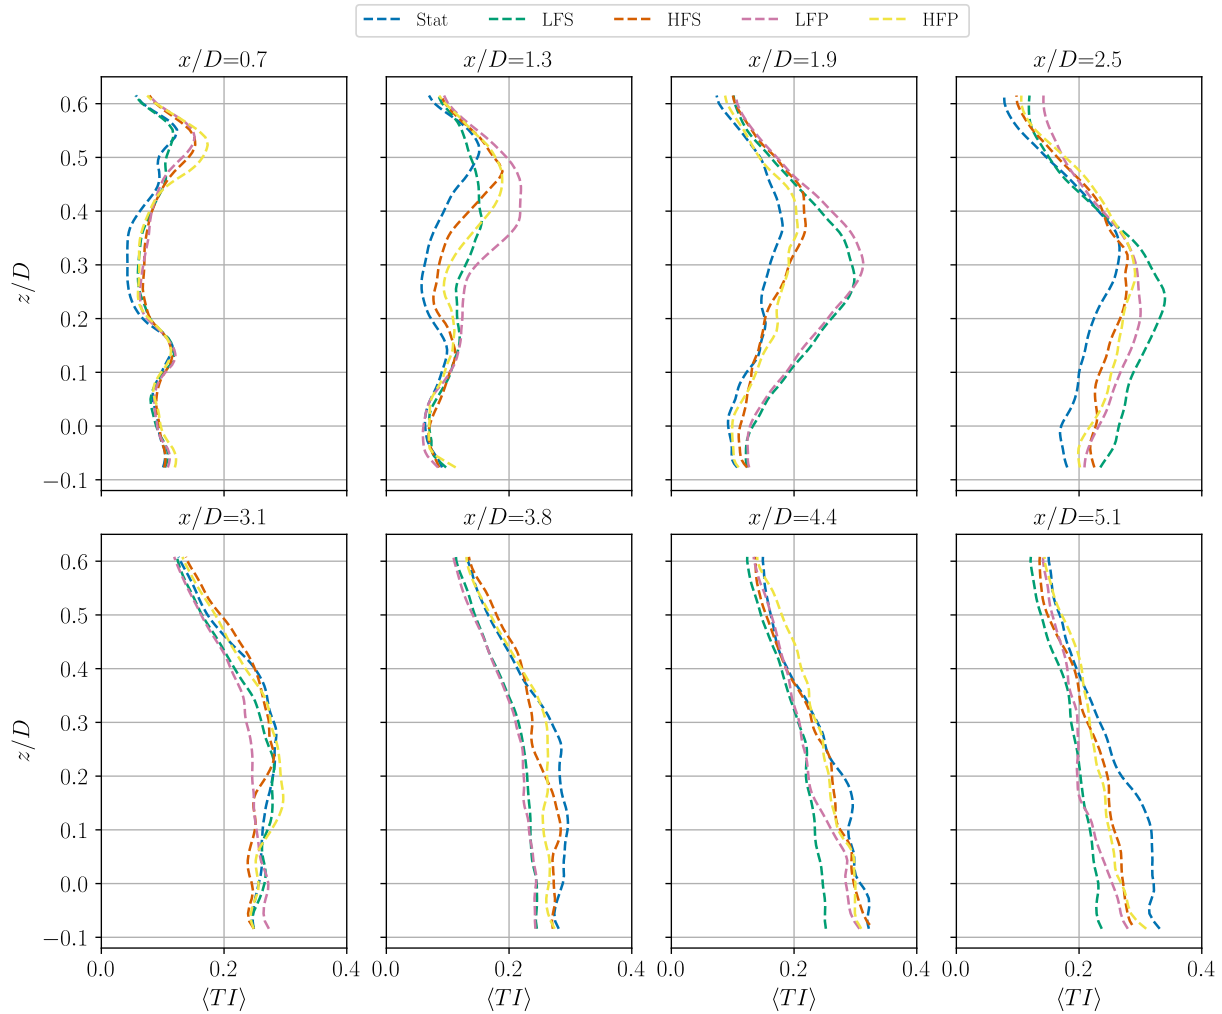


Figure 6.3: Time-averaged TI along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

6.4. Reynolds Stress

The streamwise-vertical Reynolds stress describes the turbulent flux of momentum between the streamwise and vertical directions. This is relevant for explaining why the TI evolves in a certain way to explain where momentum is extracted from outside the wake and moved into the wake. Figure 6.4 presents at eight downwind distances the streamwise-vertical Reynolds stress along the vertical direction.

6.4.1. Root Region and Wake core

At $x/D = 0.7$, within the root region, a maximum in Reynolds stress occurs at $z/D = 0.1$ and a minimum at $z/D = 0$. The minimum corresponds to turbulence transport into the lower wake core from the nacelle jet, while the maximum indicates transport into the top part of the wake. The minimum disappears further downstream near $x/D = 1.9$, likely having moved downward, below the measurement volume.

Downstream, the decreasing wake core height is apparent from the widening region where the stresses are near zero. Around $x/D = 1.9$, the positive stresses increase relatively quickly for the LFP case. Below $z/D = 0.05$, the LFS case also shows an increase, but above this height it remains the lowest among all cases. The high-frequency cases do not differ significantly from the stationary case within the root region.

Both the stationary and high-frequency cases still exhibit a region of near-zero stress, indicating the presence of a wake core. In contrast, the low-frequency cases no longer show this region. At $x/D = 2.5$, the HFP case experiences a substantial increase in Reynolds stress within the root region, while the other

cases have similar, but slightly lower, stress levels. None of the cases exhibit a region with constant zero stress here, indicating that the wake core has become fully filled with stresses and turbulence, consistent with previous observations.

6.4.2. Tip Region and Far Wake

At $x/D = 0.7$, the stationary and LFS cases exhibit relatively low stresses at this position, while the other cases have high stresses. Moving downstream to $x/D = 1.3$, the LFP case shows the largest stress beyond $x/D = 1.3$. At approximately $x/D = 1.6$, both the LFP and LFS cases display a sudden increase in stresses. The most likely reason is the breakdown of the tip vortex due to a leapfrogging instability [21]. For the stationary and high-frequency cases, this breakdown occurs later at $x/D = 2$.

Although the LFP case experiences the earliest vortex breakdown, the LFS case surpasses it after $x/D = 1.9$, reaching higher absolute stress values. After $x/D = 2.5$, stresses in the LFS case decrease, so that at $x/D = 3.1$ the stationary and high-frequency cases exhibit equally high and the highest stresses. Up to $x/D = 4.4$, the HFP case maintains the most significant stresses, followed by the stationary case.

Towards the end of the domain, the stationary case continues to show high stresses, while the HFP case exhibits a notable decrease around $z/D = 0.4$. This pattern explains the higher turbulence intensity observed for the stationary case near the wake's end.

The LFP case shows more constant and higher absolute stresses below $z/D = 0.3$ after $x/D = 3.1$, compensating for lower stresses in other wake regions and resulting in an overall TI comparable to the stationary case at the wake's end. A similar trend appears when comparing the low-frequency cases: from $x/D = 3.8$ onward, the LFP case exhibits more negative stresses near the wake axis than the LFS case, while around $z/D = 0.4$, the LFS case shows stronger negative stresses.

6.5. Vertical Velocity

The normalised time-averaged vertical velocity profiles from the experiments are presented in Figure 6.5. This metric will become relevant in Chapter 8 as a vertical position mismatch of the wake is observed between the experiments and simulations

At $x/D = 0.7$, a negative vertical velocity is observed in the nacelle jet region, with only small differences between motion cases. Above the nacelle jet, the vertical velocity gradually increases. This region, spanning from the root vortex up to the top of the wake, shows negligible variation across the different motion cases.

As the nacelle jet diffuses downstream, the vertical velocity within the jet increases for all motion cases, reaching approximately -0.03 at $x/D = 1.9$. However, above $z/D = 0.3$, the vertical velocity decreases towards negative values, and differences between motion cases start to become more pronounced—likely due to vortex breakdown. In particular, the low-frequency cases show a faster decrease than the stationary and high-frequency cases above $z/D = 0.2$.

At $x/D = 2.5$, differences between motion cases around $z/D = 0.5$ diminish again and remain small until $x/D = 3.8$. Below this vertical location, however, significant changes occur: the low-frequency cases continue to exhibit decreasing vertical velocity, while the stationary and high-frequency cases maintain positive or near-zero values below $z/D = 0.3$. For these latter cases, the vertical velocity decreases slowly downstream, with minimal profile changes up to $x/D = 3.8$.

By $x/D = 4.4$, differences between the motion cases have increased across the entire profile. The vertical velocities appear to be evolving toward uniform profiles with varying negative values. The stationary case decreases the least, followed closely by the high-frequency cases, while the low-frequency cases exhibit the lowest vertical velocity, particularly near the wake axis.

6.6. Linking the Metrics

The experimental results have been analysed for their time-averaged streamwise velocity, out-of-plane vorticity, turbulence intensity, streamwise-vertical Reynolds stresses, and vertical velocity to address

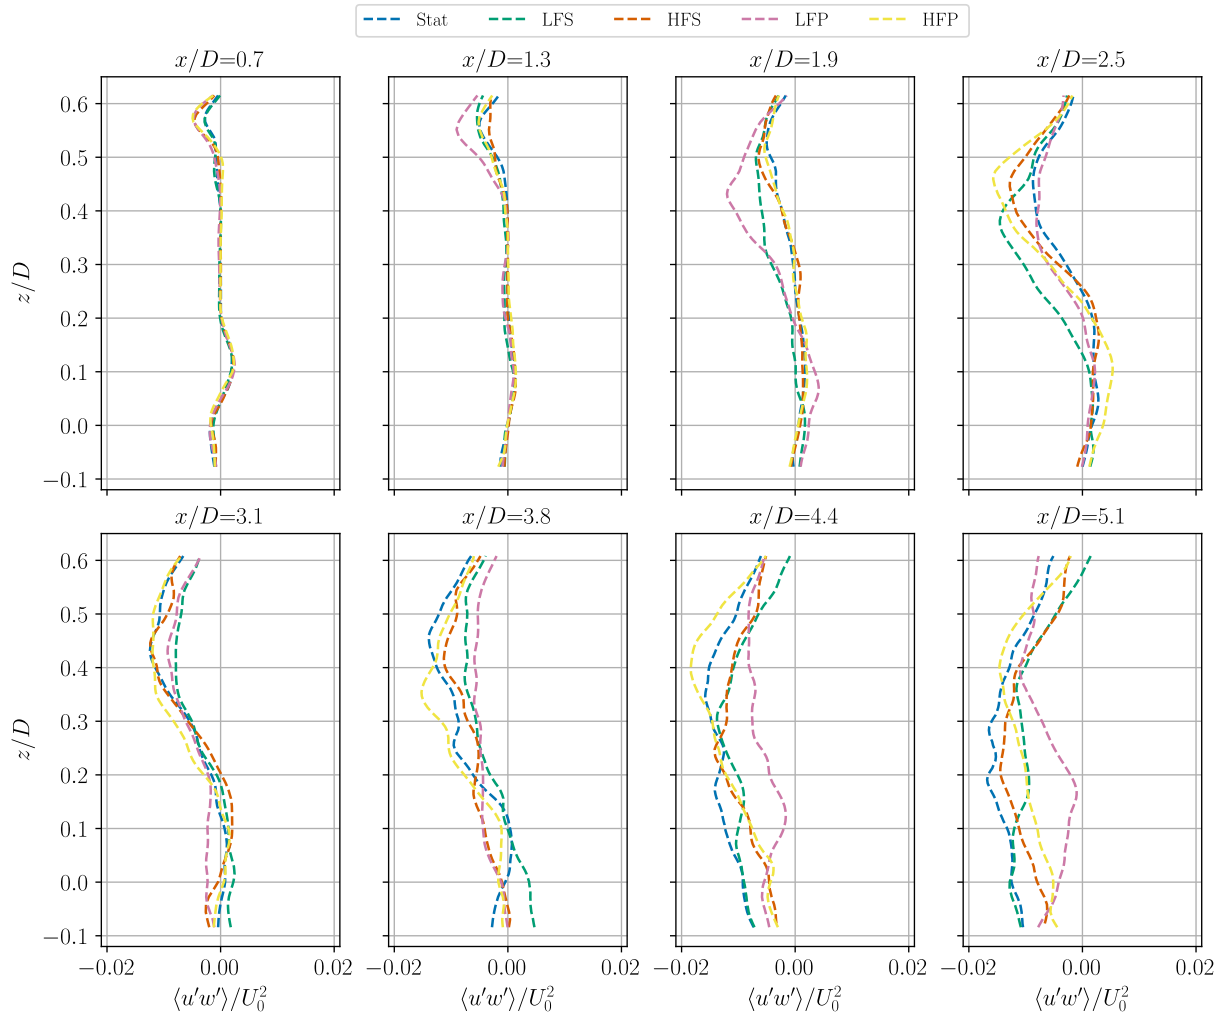


Figure 6.4: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

research question 2: “What are the characteristics of the experimentally measured wake, and what are the differences between the motion cases?” Below, the previously described observations are linked to provide a physical interpretation.

6.6.1. Root Region

The wake is initially characterized by the nacelle jet, present between $z/D = \pm 0.05$. The increased streamwise velocity in this region is explained by the out-of-plane vorticity of the root vortices. Here, a maximum vorticity from the lower half and a minimum from the upper half induce a positive streamwise velocity between these extremes. When the root vortex remains stable and concentrated, the absolute velocity gradient is high; as it breaks down and diffuses, this gradient diminishes.

The turbulence increases continuously in the regions with high vorticity since there the velocity gradients are high, causing high Reynolds stresses. The streamwise-vertical Reynolds stresses are positive above and negative below the nacelle jet, indicating turbulent transport of streamwise momentum into the wake core. The increase in turbulence causes an acceleration of the vorticity diffusion. After $x/D = 1.3$, the turbulence has led to the complete breakdown of the root vorticity below the nacelle jet, causing the nacelle and tower effects to vanish. This results in odd vorticity and symmetric streamwise velocity profiles about the wake axis. The vorticity above the nacelle jet also diffuses but persists longer, up to $x/D = 2.5$. Beyond this, the nacelle jet disappears, marking the transition from near wake to the far wake with a Gaussian streamwise velocity profile.

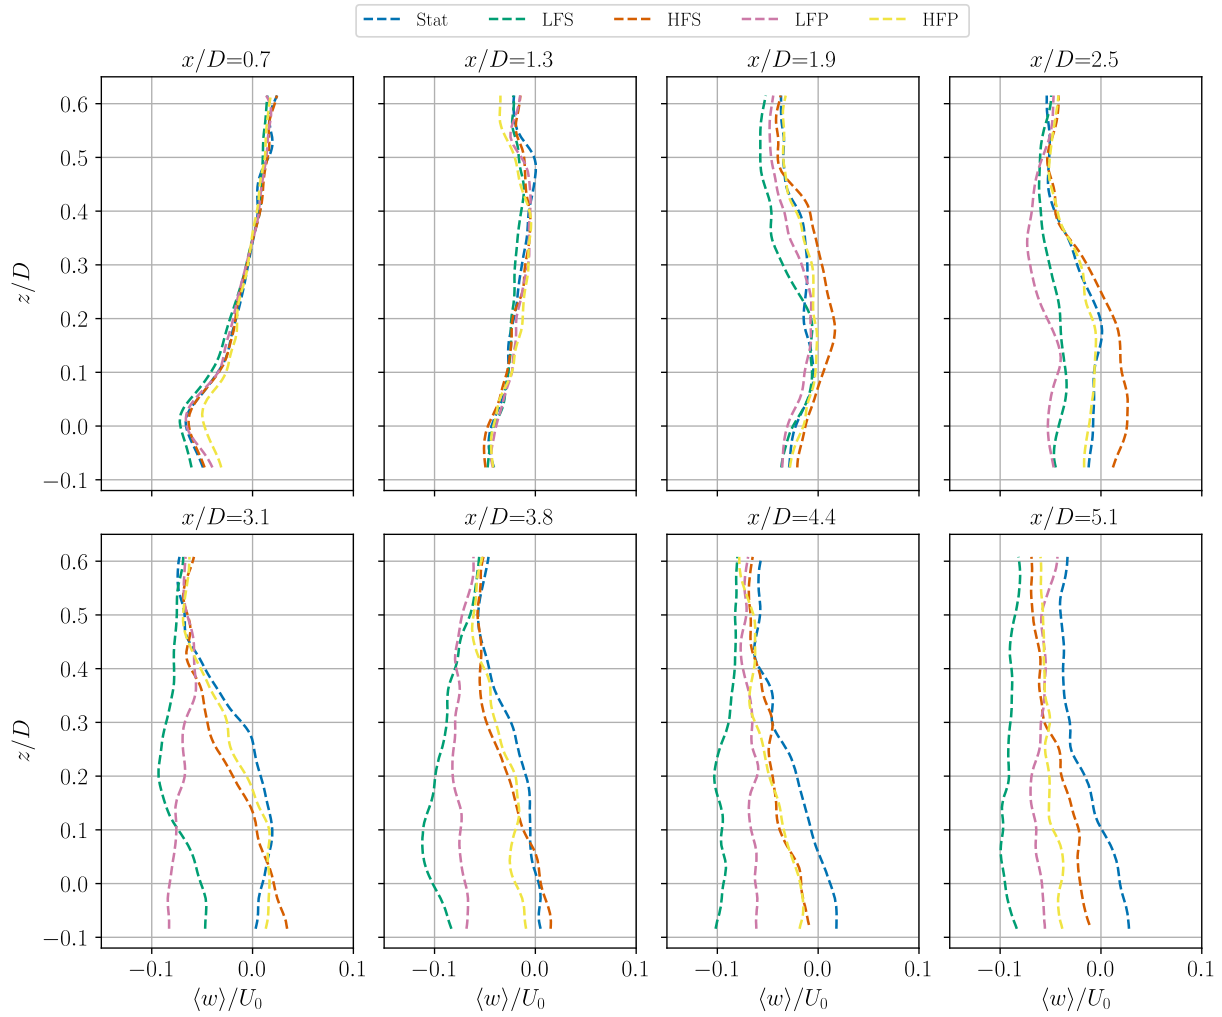


Figure 6.5: Normalised time-averaged vertical velocity along the vertical direction from the experiments for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

Differences between Motion Cases

At $x/D = 0.7$, the HFP case exhibits the highest root vorticity above the nacelle jet, slightly shifted upwards, resulting in the highest nacelle jet streamwise velocity among the cases. The low-frequency cases show higher vorticity below the nacelle jet, causing a marginally increased nacelle jet velocity at $z/D = 0$.

By $x/D = 1.9$, differences between motion cases grow, mainly due to changes in Reynolds stresses and TI. The low-frequency cases exhibit increased absolute stresses across the profile, while stationary and high-frequency cases remain unchanged. This raises the TI in the nacelle jet and accelerates vorticity diffusion, causing faster streamwise velocity decay in the low-frequency cases.

This abrupt increase in Reynolds stresses is a strong indicator of vortex breakdown, consistent with the behaviour described by [24]. The stationary and high-frequency cases show this transition later, at $x/D = 2.5$.

6.6.2. Wake Core

The wake core, situated between the root and tip vortices, initially exhibits a streamwise velocity deficit explained by the out-of-plane vorticity of these vortices, which induce velocity towards the freestream direction. At $x/D = 0.7$, turbulence in the core is minimal and originates from inflow and blade boundary layers, as evidenced by the absence of streamwise-vertical Reynolds stresses.

Downstream, turbulent mixing layers around root and tip vortices grow. At $x/D = 1.3$, Reynolds stresses

increase significantly in the tip region, with the root region following at $x/D = 1.9$. This growth raises TI levels in both tip vortex and wake core regions. Turbulence diffuses the vortices, reducing velocity gradients and narrowing the wake core.

Although vortices diffuse, the wake core velocity decreases until $x/D = 1.3$ because of a small wake expansion between $x/D = 0.7$ and 1.3 as can be seen in the vertical velocity profiles. At $x/D = 0.7$, upward velocity above the wake core indicates wake expansion. By $x/D = 1.3$, downward velocity throughout the wake contracts the wake and transports the tip vortex and turbulent stresses downward, clearly visible beyond $x/D = 1.9$.

Differences between Motion Cases

At $x/D = 0.7$, the stationary and LFS cases show an additional vortex in the tip region at $z/D = 0.47$, beneath the main tip vortex. Although the exact cause is unclear (no direct load data are available), this vortex increases the local streamwise velocity via induced velocity and enhances Reynolds stresses, resulting in a local TI maximum. Downstream, this additional vortex diffuses into the main vorticity peak but still causes a higher velocity in the LFS case.

The mixing layers, characterized by increased Reynolds stresses and TI, show growth strongly influenced by turbine motion. The LFS case experiences the fastest growth in stress magnitude and extent, while the stationary case has the lowest. At $x/D = 1.3$, HFP surpasses HFS in stresses and TI, and LFS has lower stresses than the high-frequency cases above $z/D = 0.4$. However, at $x/D = 1.9$, LFS leaps ahead, with HFS surpassing HFP.

Turbulence diffuses tip vorticity and reduces velocity gradients. Small differences in tip vortex strength are evident early, but downwind turbulence causes faster diffusion in the LFP case, followed by others, with stationary last.

6.6.3. Far Wake

At $x/D = 2.5$, a stress reduction occurs in the LFP case, followed by the HFS case around $x/D = 3.1$, while other cases maintain relatively constant stresses up to $x/D = 5.1$. This reduction coincides with the transition from near to far wake, as indicated by the Gaussian velocity profiles.

Lower Reynolds stresses cause the low-frequency cases' TI to lag behind others, with the stationary case showing the highest TI near the wake axis. These differences lead to a relatively uniform vorticity distribution in the far wake across cases. The vorticity magnitude decreases close to zero, diminishing induced velocity.

At $x/D = 2.5$, the LFP case shows the highest wake core velocity due to turbulent transport of streamwise momentum. However, the reduction in Reynolds stress slows recovery, allowing the LFS case to recover faster downstream. The HFS case also catches up with LFP due to slightly higher TI, followed by HFP and stationary cases.

Though measurements end at $x/D = 5.1$, the stationary case maintains significantly higher TI and Reynolds stresses than others, while LFP remains lowest. These differences likely impact wake recovery beyond the measurement domain; the stationary case might eventually recover faster if the measured wake were longer.

Base Simulation

This chapter discusses the results from the base configuration developed in Chapter 5. The selected metrics are discussed in Section 7.1, Section 7.3, Section 7.2, and Section 7.4, and compared against the experimental results presented in Chapter 6. Finally, the dependencies between the metrics are analysed in Section 7.5.

7.1. Streamwise Velocity

Equivalent to Section 6.1, this section presents the normalised time-averaged streamwise velocity profiles of the base configuration simulations. The profiles are along the vertical direction at eight streamwise locations in Figure 7.1. The simulation profiles are shown with solid lines, while the experimental profiles are represented by dashed lines, consistent with Chapter 6.

7.1.1. Root Region

Firstly, the nacelle jet is located between $z/D = \pm 0.05$ and is perfectly symmetric about the wake axis. The tower effect observed in the experiments is absent here, as no tower model is included in the simulations. Moreover, the streamwise velocity within the nacelle jet is overestimated due to the exclusion of the nacelle itself.

From $x/D = 0.7$ to $x/D = 1.3$, the simulated streamwise velocity decreases slightly, a minor change compared to the more significant streamwise velocity drop observed experimentally. A more pronounced change occurs at $x/D = 2.5$, where the LFP case exhibits a rapid streamwise velocity decrease in the nacelle jet, while the other cases maintain roughly constant velocities. At $x/D = 3.1$, this streamwise velocity further decreases to approximately 0.7. Around this location, other cases also begin to show small differences: notably, the LFS case has a lower jet streamwise velocity compared to the stationary and high-frequency cases.

By this point, the experiments indicate a transition toward the far wake, which the simulations do not capture well, thus inadequately predicting the near-to-far wake transition. Approaching $x/D = 3.8$, the LFP case shows an upward shift in the nacelle jet streamwise velocity peak. The other cases remain axisymmetric near the domain's end. The stationary turbine shows the second fastest streamwise velocity decrease after the LFP case. This is followed by the HFP case. At $x/D = 4.4$, the LFS and HFS cases exhibit little change; however, at $x/D = 5.1$, these also show significant streamwise velocity decreases, with the HFS case decreasing faster.

Secondly, the root streamwise velocity gradient in the simulation is located between $z/D = 0.05$ and $z/D = 0.2$ at $x/D = 0.7$, matching well with the experimental results. The increased absolute streamwise velocity observed in the experimental HFP case within the root region is not captured in the simulations.

At $x/D = 1.9$, small differences arise in the root vortex region. The LFP case exhibits a less negative streamwise velocity gradient around $z/D = 0.2$, indicating earlier recovery in the nacelle jet seen at $x/D = 2.5$, where its gradient remains the lowest among all motion cases. Other cases also display differences: the surge cases remain most stable with the highest gradients, while the HFP case shows some gradient fluctuations. Towards $x/D = 3.1$, these fluctuations also appear in the HFS case. The LFS case shows a decreasing gradient trend converging towards the LFP case, further decreasing downstream to the second lowest gradient after the LFP case. The stationary case maintains a relatively constant gradient, while the high-frequency cases show a slight gradient increase due to streamwise velocity decreases in

the wake core.

7.1.2. Wake Core and Tip Region

The streamwise velocity in the wake core is initially overestimated by 0.11, corresponding to a 30% difference. A small fluctuation around $x/D = 0.3$ is observed, which results from a low kernel size. As discussed in Section 5.4, this represents a trade-off between capturing sharp shear layers and minimizing fluctuations in the wake. Moving downstream this difference decreases.

The wake core height initially matches the experiments reasonably well, although the simulations do not capture the higher streamwise velocity seen in the stationary and LFS cases due to the additional vortex. Beyond $x/D = 0.7$, the wake core height remains nearly constant in the simulations up to $x/D = 1.9$, which contrasts with the experimental observations. Only the low-frequency cases exhibit some contraction at $x/D = 1.3$, primarily influenced by the diffusion of the tip vortex. The LFP case experiences the fastest change in wake core height, followed by the LFS case. A similar trend appears in the experiments, though the tip vortex position shifts upward, causing a vertical mismatch.

The stationary and high-frequency cases in the simulations maintain a relatively stable tip vortex and wake core, showing qualitative agreement with experiments. Approaching $x/D = 3.1$, the low-frequency cases display a further decrease in the tip vortex streamwise velocity gradient, while other cases remain relatively constant. The stationary case shows a lower tip vortex gradient than the high-frequency cases, opposite to what is observed experimentally.

Within the wake core, the low-frequency cases continue to see a decrease in streamwise velocity due to tip vortex diffusion, although this decrease is insufficient to match experimental results. At $x/D = 3.8$, the LFP case exhibits an increase in core streamwise velocity associated with vortex structure breakdown. The LFS case also shows a similar velocity increase, consistent with experimental observations of vortex breakdown around $x/D = 1.6$.

Above $z/D = 0.3$, the streamwise velocity of the LFS and LFP cases are comparable, but below this height, the LFP case demonstrates a stronger speed-up, as also seen in the experiments at $x/D = 2.5$. Further downstream, the LFS case approaches the stationary and high-frequency cases above $z/D = 0.3$. However, the experiments reveal a significantly higher recovery rate, resulting in a substantial underprediction of the streamwise velocity in the simulations.

7.2. Vorticity

This section presents the normalised time-averaged out-of-plane vorticity profiles of the base configuration simulations. The profiles are along the vertical direction at eight streamwise locations in Figure 7.2. Simulation results are shown as solid lines, while experimental data are shown as dashed lines (as in Chapter 6).

7.2.1. Root Region

In the nacelle jet, the simulations predict zero vorticity throughout the domain. Flanking this region are two peaks corresponding to the root vortices. Their profiles appear asymmetric about the wake axis due to the absence of the tower and nacelle in the model, which is in contrast to the experiments. The minimum, located at $z/D = 0.12$, aligns with the experimental position and coincides with the centre of the root vortex gradient identified in Section 7.1. The simulated vortex width is well captured, but its peak vorticity is overestimated by approximately 20%. This discrepancy increases downstream. While the simulations show growing vorticity up to $x/D = 1.9$, the experiments exhibit a decrease after $x/D = 1.6$ due to vorticity diffusion. At $x/D = 1.9$, the LFP case shows a slight reduction, followed by a further drop at $x/D = 2.5$; a similar trend is seen for the LFS case.

Additionally, the root vorticity of all cases shows the formation of a secondary vortex beyond $x/D = 1.9$ at $z/D = 0.15$. This could be the result of mutual-inductance of two vortices causing them to rotate around each other. However, it is expected that this would lead to root vortex breakdown which is not observed. Anyway, this additional vortex is not observed in the experiments.

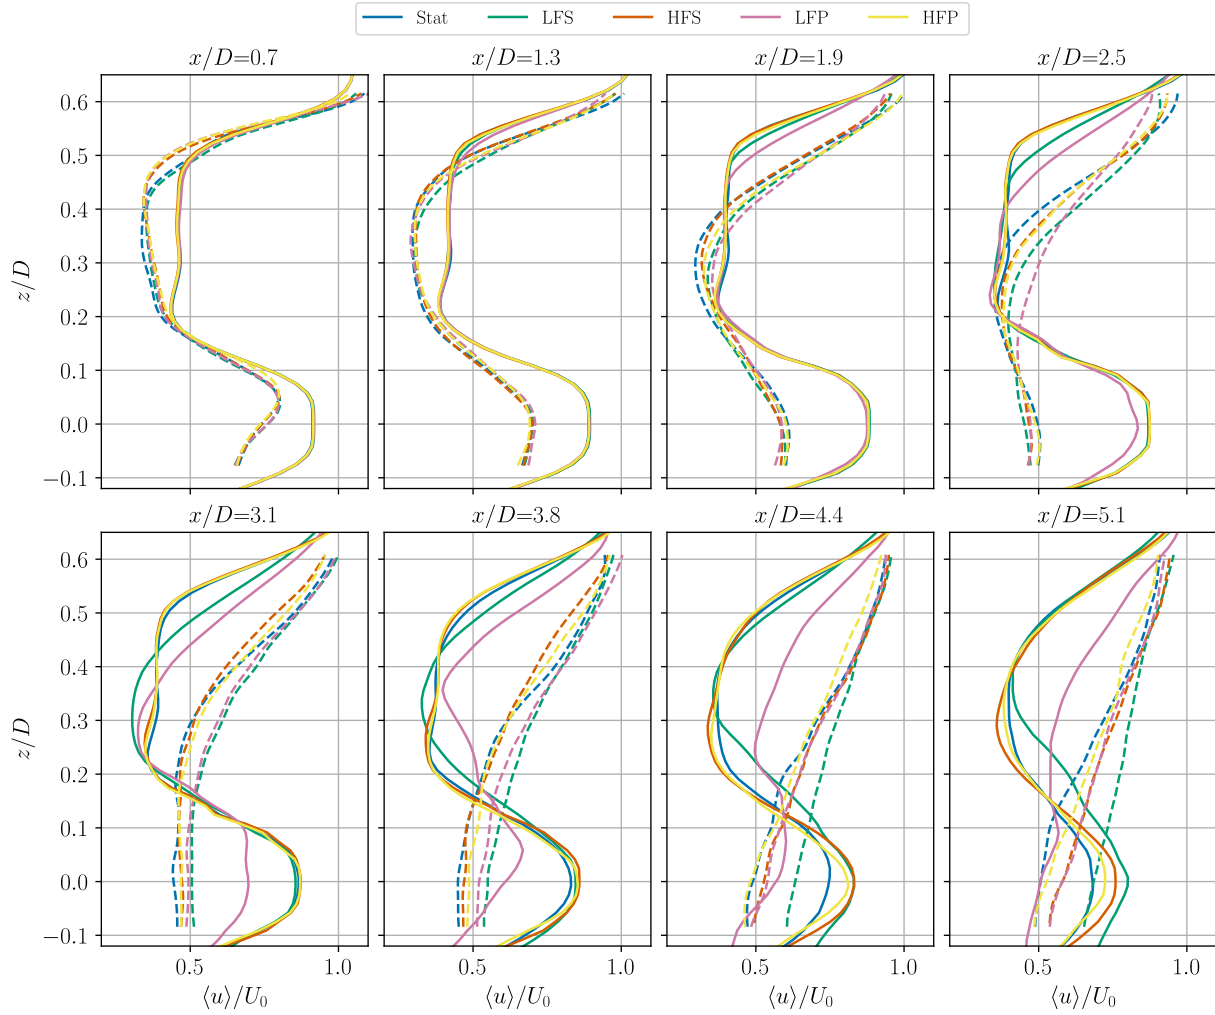


Figure 7.1: Normalised time-averaged streamwise velocity along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

At $x/D = 2.5$, the LFP case shows an upward shift of the root vorticity, indicating a possible instability. The other cases and experiments do not show this.

The low-frequency cases consistently have the weakest vortices, which is also observed in the experiments. Between the high-frequency cases and the stationary case there is little difference until $x/D = 3.1$, where the high-frequency cases have more concentrated root vorticity.

7.2.2. Wake Core and Tip Region

The wake core initially has near-zero vorticity, reflecting the low blade load gradients in the mid-span and the roll-up of vorticity into the root and tip vortices. Above $z/D = 0.2$, a slight vorticity increase is observed, likely causing the small velocity fluctuation in the wake core through its induced velocity. Above this point, vorticity returns to zero, until the tip vortex produces a local maximum.

In contrast to the experimental results, which show significant variation in tip vortex strength, the simulations predict no such variation. The tip vortex strength is also underestimated, with the stationary turbine showing the largest error (33%) and the LFS case the smallest (15%). Furthermore, the secondary vorticity maximum at $x/D = 0.7$ seen just below the main peak in the stationary and LFS experimental cases is absent in the simulations.

Downstream, the wake core expands due to the upward motion of the tip vortex. This expansion continues until $x/D = 1.5$ for the low-frequency cases and until $x/D = 3.1$ for the stationary and high-frequency

cases, whereas the experiments show expansion ending before $x/D = 0.7$. As the tip vortex convects downstream, viscous diffusion broadens and weakens the maximum. At $x/D = 1.3$, the relative strength between cases is captured: the stationary case and high-frequency cases are the strongest, followed by the LFS, while the LFP remains weakest in both simulation and experiment. However, the simulations overpredict absolute vorticity due to underpredicted vortex diffusion until the end of the measurement domain. Moreover, by $x/D = 2.5$, experiments show the tip vortex maximum moving toward the wake axis, but the simulations do not capture this.

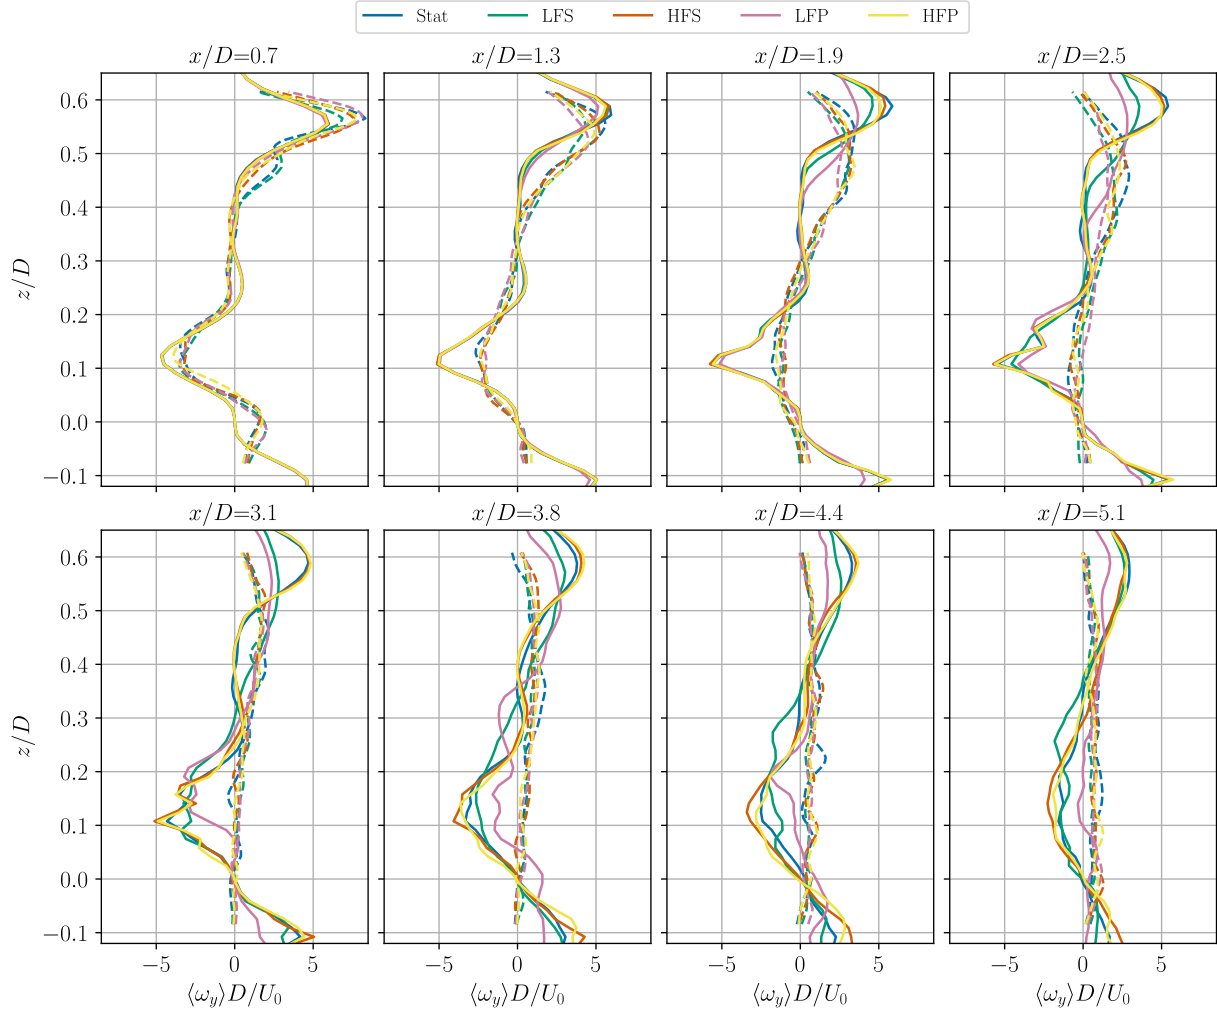


Figure 7.2: Normalised time-averaged out-of-plane vorticity along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

7.3. Turbulence Intensity

Figure 7.3 presents the time-averaged TI for the base configuration simulations. The profiles are along the vertical direction at eight streamwise locations. Simulation results are shown as solid lines, while experimental data are shown as dashed lines (as in Chapter 6).

7.3.1. Root Region

In the root region, TI is lower in the simulations than in the experiments. This is due to excluding any inflow turbulence and due to the absence of the nacelle and tower. At $x/D = 0.7$, the LFP case exhibits a small increase at $z/D = 0.12$ associated with the root vortex. By $x/D = 1.3$, the LFP case has increased several percent, which is more than observed in the experiments.

At $x/D = 1.9$, TI in the nacelle jet increases by 1% for all cases except LFP, which rises to 6%. The LFS case shows a sudden increase in the root vortex, followed by an even larger increase at $x/D = 2.5$. At $x/D = 3.1$, the stationary case and high-frequency cases show early signs of turbulence in the root vortex at $z/D = 0.12$. Also, the differences between the motion cases in the nacelle jet increase. Moreover, the LFP profile shifts upward, consistent with earlier velocity and vorticity observations.

At $x/D = 3.8$, the stationary case catches up to LFS in the nacelle jet, surpassing it by $x/D = 3.8$. Also the high-frequency cases surpass it by $x/D = 5.1$. This acceleration of turbulence production is also observed in the experiments, however, already after $x/D = 2.5$. Next to this, the LFP case remains at a high TI relative to the other cases in contrast to the experiments.

7.3.2. Wake Core and Tip Region

At $x/D = 0.7$, simulations show increased TI in the wake core compared to the root region due to turbulence shedding of the blades. The increase is more pronounced for moving turbine cases than for the stationary turbine, consistent with the experiments. In both simulation and experiment, the stationary turbine has the lowest TI, though the ranking of the moving cases differs from experiments. Simulated TI values are generally lower. By $x/D = 1.3$, TI increases and the order of high to low matches the experiments, but with smaller differences between cases.

The turbulence in the tip region is increasing faster than that of the wake core, resulting in a turbulence minimum in the wake core at $x/D = 1.9$. This process is fastest for the low-frequency cases. At this point, a similar minimum has already disappeared in the experiments due to the growing mixing layer in the tip region.

At $x/D = 0.7$, there is a small TI maximum in the simulations, located at $z/D = 0.55$ for low-frequency and $z/D = 0.48$ for high-frequency cases; the reason for this shift is unclear. In the experiments there are maxima too, however, they are 9–15% higher. By $x/D = 1.9$, the low-frequency cases show a sudden turbulence increase, which persists until $x/D = 3.1$ where the TI becomes relatively constant. For the stationary case and the high-frequency cases this increase in turbulence starts accelerating around $x/D = 3.1$ while in the experiments a similar event is observed around $x/D = 1.9$.

Beyond $x/D = 3.8$, LFS has higher TI in the tip region than LFP, and the stationary case surpasses the high-frequency cases, a trend also seen experimentally but further downstream. At $x/D = 4.4$, simulated LFS TI decreases slightly, aligning more closely with experiments but still overpredicting. The stationary and high-frequency cases continue to increase, with the maximum at $z/D = 0.5$ matching the experimental peak location but with significantly lower values below it. By $x/D = 5.1$, tip-vortex-induced TI in stationary and high-frequency simulations is overpredicted at $z/D = 0.5$, while most cases still underestimate TI below this height.

7.4. Reynolds stress

Figure 7.4 presents at eight downwind distances the streamwise-vertical Reynolds stress for the base configuration simulations. The profiles are along the vertical direction at eight streamwise locations. Simulation results are shown as solid lines, while experimental data are shown as dashed lines (as in Chapter 6).

7.4.1. Root Region

Initially, stresses are effectively zero across the vertical direction for all cases. At $z/D = 0.1$ and on the wake axis, the LFP case shows a small maximum, negligible compared with the experimental maximum at the same height.

By $x/D = 1.9$, the LFP stress at $z/D = 0$ becomes more negative, opposite to the experimental trend. This indicates momentum transfer from the wake core downward, counteracting wake recovery. Other cases remain near zero while experiments show increasing stress. The stress of LFP case continues decreasing until $x/D = 3.1$, coincidentally matching the experimental value when the measured wake has already transitioned to the far wake.

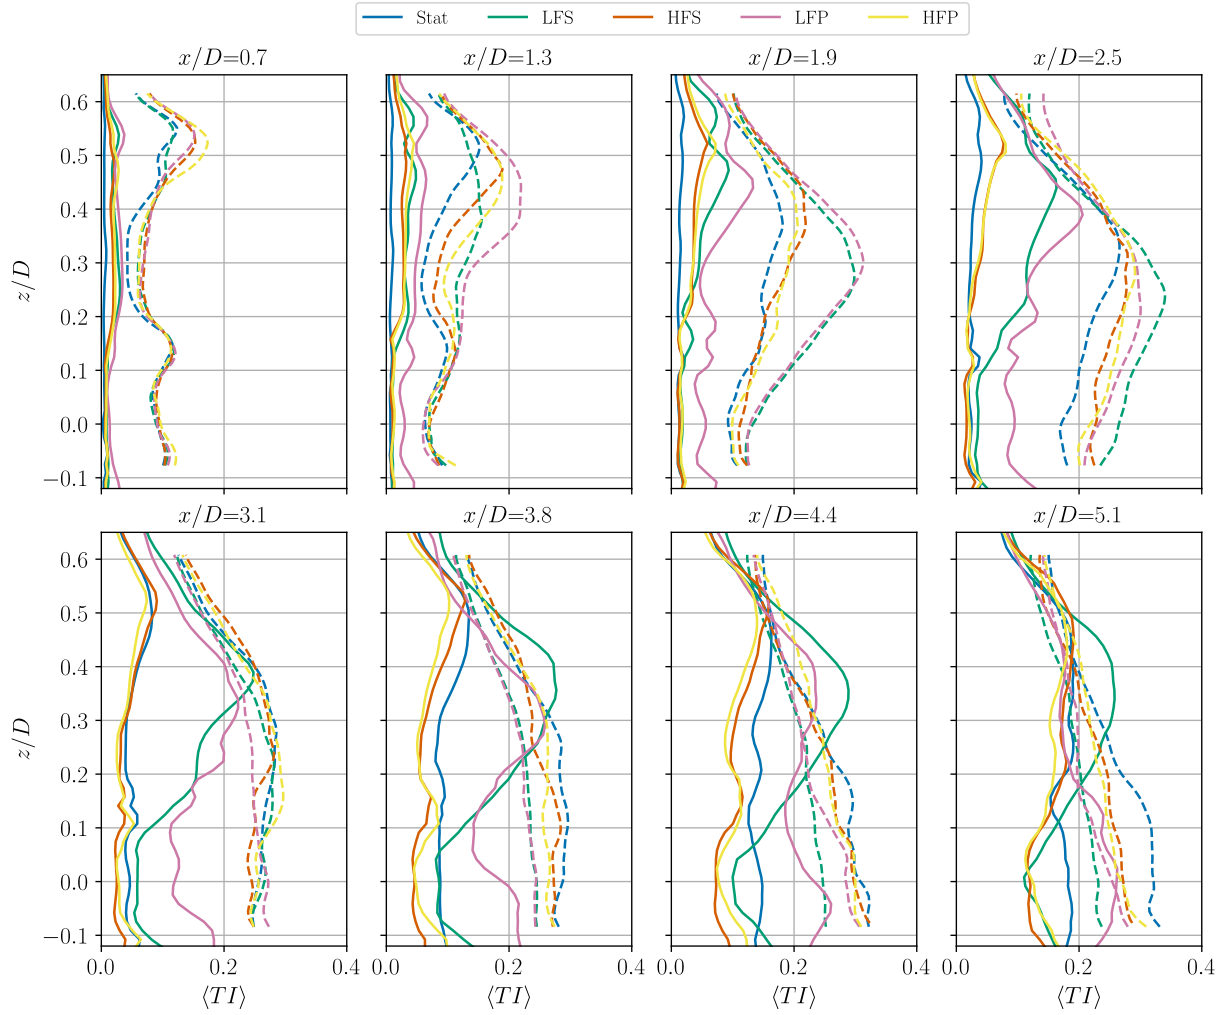


Figure 7.3: Time-averaged TI along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

At $x/D = 3.1$, around $z/D = 0.1$, stationary, LFS, and high-frequency cases exhibit small stress increases. By $x/D = 3.8$, LFP shows a large stress increase in the root region and above, whereas experiments show a decreasing stress in the root region. However, towards $x/D = 5.1$, LFP returns to negative stress, and the other cases fluctuate around the axis (negative below, positive above), a general behaviour very different from experiments.

7.4.2. Wake Core and Tip Region

In the wake core at $x/D = 0.7$, simulations correctly show near-zero stress, though the minimum in the tip region is barely visible. At $x/D = 1.3$, the wake core remains well captured, but experiments display small deviations from zero absent in the simulations. Simulations show a shallow minimum in the tip region, positioned too high due to the previously noted increased wake expansion. In both simulation and experiment, LFP exhibits the largest absolute stress, followed by LFS and HFP; HFS shows a weak minimum, and the stationary case incorrectly shows none.

At $x/D = 1.9$, low-frequency cases still have the highest stress, underpredicted relative to experiments and positioned too high. The simulated case ranking is LFP highest, then LFS, then high-frequency, and stationary lowest. This matches the experiments. The width of the minimum increases too as a result of the mixing layer growth. However, it remains behind that of the experiments.

By $x/D = 2.5$, larger fluctuations in the LFP case can be observed when travelling downstream. As a

result there is a very positive stress in the wake core at $x/D = 4.4$, while in the experiments the stress is negative. The other cases evolve more consistently but at a slower pace to negative stresses in the wake core and tip region but generally underpredict the stress. Consistencies in differences between motion cases are difficult to identify in this part of the wake. It can be noted that the stationary case has the most significant stress underprediction, followed by the high-frequency cases. The profiles correspond better with the more upwind profiles.

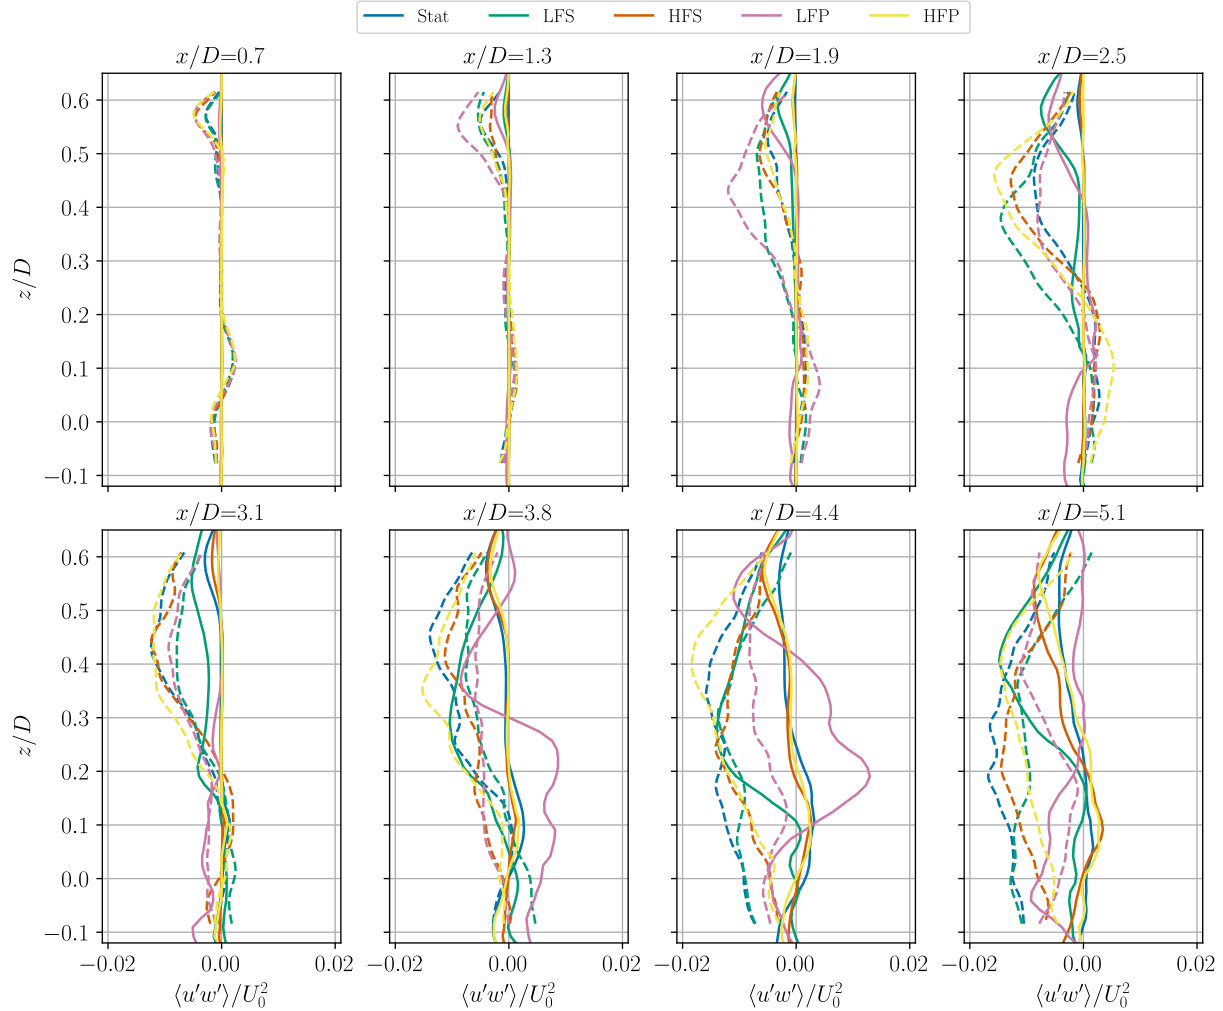


Figure 7.4: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the base simulation (full line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

7.5. Linking the Metrics

This section examines how the measured metrics influence each other to provide a complete understanding of the wake physics, enabling proposals for simulation improvements. The wake is divided into the root region, including the nacelle jet and root vortices, and the wake core, which is strongly influenced by both root and tip vortices.

7.5.1. Root Region

Immediately after the rotor plane, the nacelle jet streamwise velocity is overpredicted and symmetric about the wake axis. This arises from the exclusion of tower and nacelle effects, which would otherwise break the axisymmetry and introduce turbulence. The symmetry itself comes from odd root vorticity, producing equal induced velocity from the upper and lower root vortices. Overprediction of root vorticity further enhances induced velocity, amplifying the nacelle jet overprediction. Including tower and nacelle

effects could significantly improve the near wake up to $x/D = 1.3$.

In experiments, nacelle jet vorticity diffuses much earlier than in simulations, by $x/D = 1.9$ in the former and $x/D = 2.5$ in the latter. The delay is linked to negligible Reynolds stresses and thus a low TI until $x/D = 1.9$, when LFP develops a small increase. At $x/D = 2.5$, LFP experiences non-negligible stress, increasing TI and enabling stronger vorticity diffusion. The other cases develop at an even lower rate.

Stationary and LFS cases develop the second-highest stresses, producing a gradual TI rise and further nacelle jet decay. At $x/D = 3.8$, LFS stress growth halts, allowing high-frequency cases to catch up in TI and vorticity diffusion. The LFS case exhibits the widest stress and TI footprint, producing broader vorticity diffusion and induced velocity into the wake core. However, its nacelle jet velocity at $x/D = 5.1$ remains the highest of all cases due to persistent velocity induction.

7.5.2. Wake Core

Throughout the measured domain, the wake core streamwise velocity is overpredicted. This results from underpredicted tip vorticity, which reduces induced velocity. While the tip vorticity is low, it is sufficient to match the observed streamwise velocity gradient magnitude. Slightly overpredicted root vorticity steepens the root-vortex gradient, but combined induced velocity at $z/D = 0.2$ is accurate at $x/D = 0.7$.

The tip-vorticity deficit increases momentum in the core, sustaining overprediction downstream. This shortfall may stem from insufficient tip downwash, potentially correctable by the FALM or by refining the mesh to reduce numerical vorticity diffusion and improve tip-force gradients.

Small fluctuations in core velocity arise from too low a kernel size, which imprints on the vorticity field. These fluctuations persist to $x/D = 4.4$ in stationary and high-frequency cases, and to $x/D = 2.5$ in low-frequency cases. Increasing kernel size could suppress them but would degrade tip-region velocity gradients, as discussed in Section 5.4.

At $x/D = 1.9$, simulated tip vortices rise, indicating wake expansion absent in experiments. This shifts the tip-region velocity gradient vertically, although magnitude remains correct. Expansion causes a core velocity drop for continuity, persisting while expansion continues.

Tip-vortex diffusion is turbulence-driven, but underpredicted shear stresses reduce turbulence generation, prolonging vortex stability and delaying recovery. LFP has the highest stresses to $x/D = 1.9$, followed by LFS; after $x/D = 2.5$, these two cases have similar stresses. Other cases remain near zero stress and low TI, slowing tip-vorticity decay. In low-frequency cases, tip-vortex diffusion reduces velocity gradients as in experiments, but later. Stationary and high-frequency cases only show this decrease by $x/D = 3.8$, when TI has risen sufficiently; the stationary case exhibits slightly greater diffusion than high-frequency cases.

The low stress and TI levels result from low-turbulent inflow. Adding realistic inflow TI would promote earlier vorticity breakdown, as also observed in [59].

LFP

While most cases recover similarly, the LFP case recovers fastest due to a consistent upward swept root vorticity, increasing stresses and TI near the root vortex. This upward push accelerates diffusion and redistributes vorticity such that induced velocity is reduced relative to other cases. The force likely originates from neighbouring root vortices rotating around each other, as seen in vorticity profiles, leading to mutual destabilisation. Although other cases show similar rotation, they remain stable, possibly due to weaker vorticity magnitude and vortex strength.

Turbulent Inflow

As shown in Chapter 7, the laminar inflow simulations produced an overly stable wake characterised by an excessively long expansion region, low TI, and limited wake entrainment. The primary cause is the low inflow turbulence. This chapter examines the wake behaviour of the simulated motion cases when subjected to turbulent inflow.

First, Section 8.1 describes the method used to generate the turbulent inflow and summarises its key characteristics. The wake development is then analysed in terms of streamwise velocity (Section 8.2), out-of-plane vorticity (Section 8.3), TI (Section 8.4), and streamwise–vertical Reynolds stress (Section 8.5). For each of these metrics, the wake is divided into four parts: the root region, wake core, tip region and far wake. In addition to the regular metrics, Section 8.6 discusses the vertical velocity field, which helps explain specific discrepancies observed within the wake. Finally, the interdependencies between these metrics are considered in a combined analysis in Section 8.7.

8.1. Turbulence Generation

In GRASP, turbulence is generated by applying a heat and/or humidity flux at the bottom boundary, as described in Section 8.1. In this study, only a heat flux is applied. The turbulence intensity in the domain is controlled by two parameters: the perturbation magnitude and the heat flux. As discussed in Section 5.8, a stable TI is obtained with a small perturbation magnitude.

In the wind tunnel, the TI at the turbine location was approximately 3%, slightly higher than the minimum recorded by [21]. This may be due to the turbine’s off-centre position in the wind-tunnel jet, causing part of the rotor to pass through the turbulent shear layer.

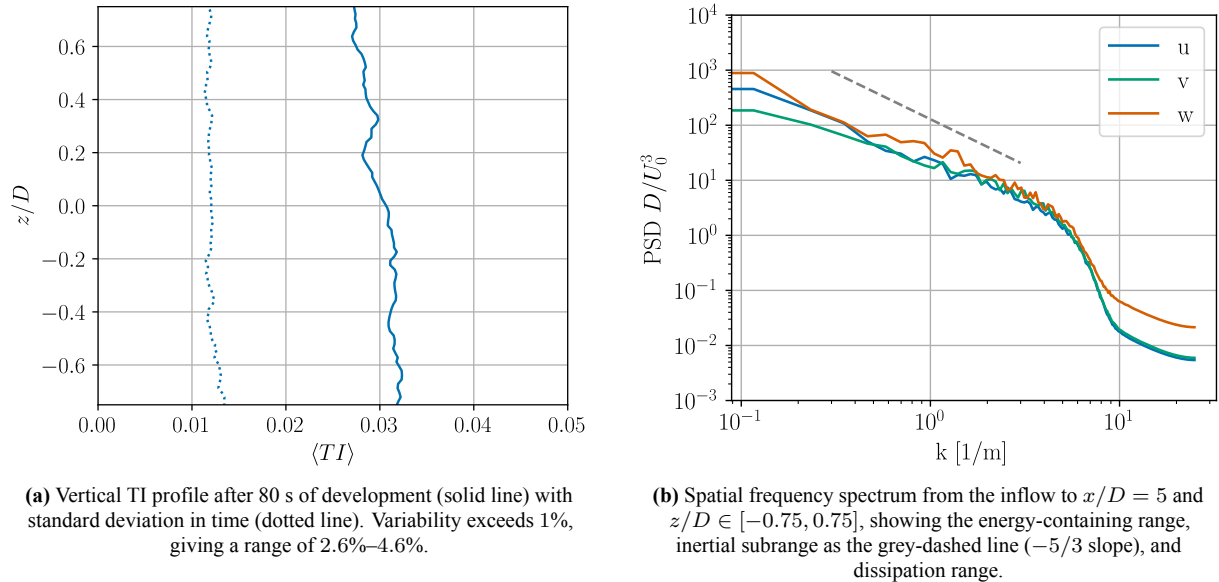


Figure 8.1: Turbulent inflow characteristics.

8.1.1. Inflow TI Profile

A heat flux of 0.1 Km/s produced the TI shown in Figure 5.14b, which was higher than the target value. Therefore, a reduced flux of 0.08 Km/s was selected. The precursor simulation was run for 80 s to allow

the inflow to develop to a uniform TI of 3.6%.

Figure 8.1a shows the vertical TI profile (solid line) averaged over 10 s, together with its standard deviation in time (dotted line). The standard deviation exceeds 1%, meaning most values lie between 2.6% and 4.6%. Attempts to reduce this variability were unsuccessful.

8.1.2. Turbulence Spectra

The development of turbulence in the precursor simulation was assessed from the time and spatial spectra of the streamwise, vertical, and lateral velocities. Both spectra showed similar behaviour; only the spatial spectrum is shown in Figure 8.1b.

The spectra display all features of well-developed turbulence: an energy-containing range at low wavenumbers, a $-5/3$ slope in the inertial subrange (dashed grey line), and a dissipation range at high wavenumbers. The streamwise and vertical components contain more energy in the large scales than the lateral component, due to the streamwise mean flow and the bottom-boundary heat flux. Despite these differences, the cascades are very similar, indicating isotropic turbulence in the turbine inflow.

8.1.3. Sub-Grid Scale Model Activity

Adding turbulence to the inflow alters the resolved-to-total TKE ratio. For laminar inflows, Section 5.5 showed that Pope's criterion could not be satisfied because the SGS TKE contribution was high in the tip and root vortices, requiring very fine resolution.

With turbulent inflow, large eddies carry significant resolved TKE into the domain, reducing the relative SGS contribution. For the stationary case, Figure 8.2 shows that the resolved TKE fraction exceeds 90% everywhere in the wake, easily satisfying Pope's criterion. This confirms that the chosen mesh resolution is adequate, and highlights that the criterion is more challenging to meet for laminar inflows where turbulence is generated only by boundary-layer and wake shear.




Figure 8.2: Resolved TKE as a fraction of total TKE for the stationary turbine with a uniform 3.6% TI inflow.

8.2. Streamwise Velocity

Normalised time-averaged vertical velocity profiles at eight streamwise locations are shown in Figure 8.3. Simulation results are plotted as solid lines, experimental data as dashed lines.

8.2.1. Root Region

At $x/D = 0.7$, the nacelle jet streamwise velocity is still overpredicted because no nacelle or tower is modelled. Compared to the low-turbulence inflow, the transition from the jet maximum to the root vortex gradient at $z/D = \pm 0.05$ is smoother. The root vortex gradient is reproduced accurately, with no visible differences between motion cases in the simulations, whereas the experiments show small variations.

At $x/D = 1.3$, the nacelle jet velocity decreases in both simulations and experiments. In the experiments, the maximum drops rapidly, producing a mismatch in streamwise velocity within the root vortex region. Nevertheless, simulated gradients match the experimental ones. Experimental case-to-case variations remain absent in the simulations.

By $x/D = 1.9$, the nacelle jet velocity is closer to experimental values, with a relative error of 4%. The absolute root vortex gradient increases at the same rate in both datasets. Some motion case variation appears in the nacelle jet: the LFP case has the highest simulated velocity but is lowest in the experiments; LFS shows the opposite trend. Variations in the root vortex are similarly inconsistent and negligible.

At $x/D = 3.1$, the nacelle jet is nearly diffused, consistent with the experiments, though still slightly overestimated. Experimental profiles show large motion case differences, but simulations only fluctuate slightly around each other. The LFS case consistently produces a lower velocity just above the wake axis in both datasets. At this point, the tip vortex has dissipated and moved downward, leaving little of the root vortex gradient.

8.2.2. Wake Core

At $x/D = 0.7$, the wake core velocity is the same as for the base simulations. However, its width has been reduced slightly since the tip velocity gradient magnitude has decreased due the turbulence. Also, the fluctuation at $z/D = 0.25$ has become smaller.

By $x/D = 1.3$, the velocity decreases and the wake core narrows due to a decrease of the root and tip velocity gradient magnitudes. Although, the overprediction is still visible in the wake core, the general shape of the profiles are consistent with the experiments.

At $x/D = 1.9$, the wake core is more rounded, showing a single velocity maximum rather than a flat vertical profile, though the absolute velocity is still overpredicted. Also, the velocity minimum starts to increase, indicating wake recovery. Motion case differences remain small but consistent compared to the experiments: low-frequency cases show higher velocities than the others. These differences increase toward $x/D = 2.5$. A slight velocity increase for the high-frequency cases relative to the stationary case also appears, matching experimental trends. The simulated wake core position, however, does not align with the measurements as the experiment profiles are shifted downward.

8.2.3. Tip Region

At $x/D = 0.7$, the tip vortex produces a mean velocity gradient ending near $z/D = 0.6$ at a normalised velocity of 1.05. This differs from the experiments, where the gradient is steeper and continues to increase beyond 1.09. In the experiments, the LFP and high-frequency cases have higher gradients, and the stationary and LFS cases probably as well if additional vorticity were not present below $z/D = 0.5$.

At $x/D = 1.3$, simulations and experiments show a similar decrease in gradient, except in the stationary case, where the tip region maintains a constant velocity. Minor motion case variations appear in the simulations: the LFP case has the lowest velocity gradient, followed by the LFS case, though in reverse order to the experiments. In both datasets, these two cases remain lower than the stationary and high-frequency cases. At $z/D = 0.6$, all cases asymptotically approach a normalised velocity of 1, at a slightly lower position than in the experiments due to the weaker gradient.

At $x/D = 1.9$, the simulated gradients remain similar to the experimental values. However, the vertical shift causes a mismatch. Motion case differences are small and similar to what was observed previously. By $x/D = 2.5$, the velocity is overpredicted in the tip region, and the simulation's ordering of motion case variations no longer matches the experiments.

8.2.4. Far Wake

The far wake begins around $x/D = 3.1$, where the velocity profiles start to exhibit a Gaussian shape as the nacelle jet vanishes. Significant differences between the motion cases are observed in this region. In the simulations, the velocity profiles fluctuate around each other without a clear trend, unlike the consistent variations seen in the experiments.

Up to approximately $z/D = 0.3$, the simulated velocity profiles closely match the experimental shape, with a normalised velocity near 0.5 at the wake axis. Above this height, the simulations tend to underpredict the velocity, as was also the case at $x/D = 2.5$. This could be due to the observed downward shift of the profiles. This underprediction remains until the end of the domain.

At $x/D = 3.8$, the variation between motion cases in the simulations resembles the experiments: the low-frequency cases exhibit higher wake-axis velocities compared to the stationary and high-frequency cases. Above $x/D = 0.2$, the LFS case has the lowest velocity and the LFP case the highest, followed by the HFP case. The other cases do not show significant differences. By $x/D = 4.4$, the experimental differences between motion cases increase across the full vertical extent, but the simulations fail to capture this growth.

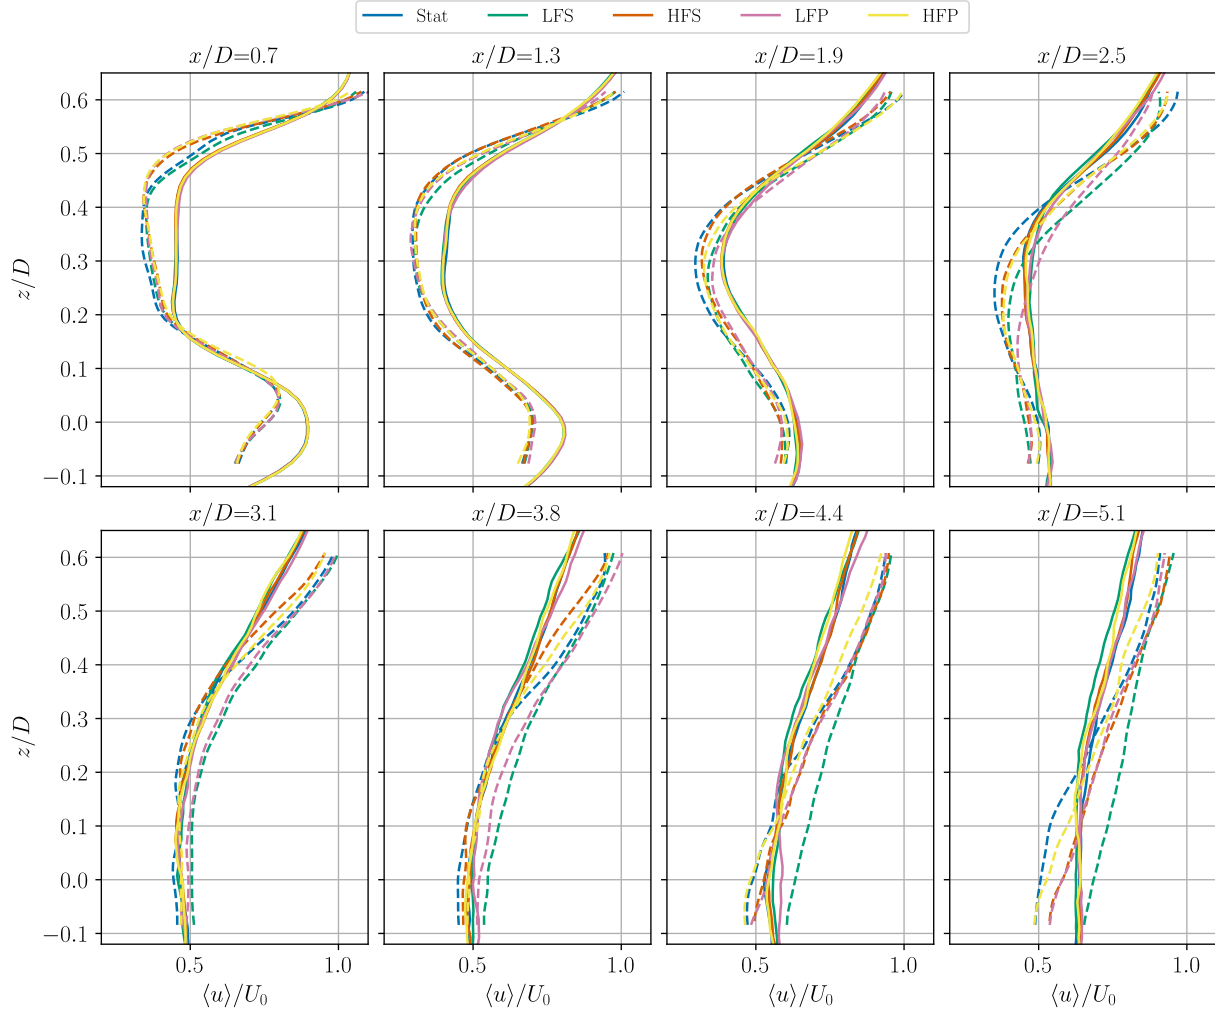


Figure 8.3: Normalised time-averaged streamwise velocity along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

8.3. Vorticity

This section presents the normalised time-averaged out-of-plane vorticity profiles at eight streamwise locations as shown in Figure 8.4. Simulation results are plotted as solid lines, experimental data as dashed lines.

8.3.1. Root Region

At $x/D = 0.7$, the vorticity is zero at the wake axis. Around this axis a maximum and a minimum are visible, corresponding to the root vortices. The magnitude of the minimum is predicted very well by the simulation, whereas the base simulations overpredicted the absolute vorticity. As in the velocity results, differences between motion cases are not captured. The vorticity maximum is at a lower vertical position in the simulations owing to the omission of the tower and nacelle. The location of the minimum is also lower in the simulations, likely a consequence of that omission.

By $x/D = 1.3$, the root vortex has diffused significantly. The upper-root vorticity distribution is well reproduced, although there remains no variation between motion cases. Below $z/D = 0.05$, a small underprediction of the absolute vorticity is observed, again probably caused by the missing tower and nacelle.

From this point until the end of the near wake at $x/D = 2.5$, the root vortex gradually diffuses until it is effectively zero. This diffusion is captured accurately by the simulations, even though the differences between motion cases remain effectively absent.

8.3.2. Wake Core

At $x/D = 0.7$, the core height is slightly larger in the simulations due to the downward shift of the root vortex. Moreover, compared to the low-turbulence inflow results, the small maximum at $z/D = 0.25$ due to the too low kernel size has decreased significantly due to the turbulence.

At $x/D = 1.3$, the wake core extends accurately from $z/D = 0.3$ to $z/D = 0.38$. Negligible differences between motion cases are simulated. However, these are also small in the experiments.

At $x/D = 1.9$, there is effectively no wake core remaining, as the root and tip vorticity have completely diffused towards each other at $z/D = 0.3$. At this height, the slope of the vorticity matches very well with the experiments.

8.3.3. Tip Region

At $x/D = 0.7$, the out-of-plane tip vorticity is at maximum underpredicted by 44%. The underprediction becomes less downstream but never fully disappears. Additionally, the significant variation between the motion cases is not captured, including the local vorticity increase in the stationary and LFS cases. Moreover, the position of the maximum is lower than in the experiments.

At $x/D = 1.3$, the LFP case shows a smaller vorticity at the maximum, which could indicate the effect of the motion, also observed in the experiments. Towards the edges of the tip vorticity, both low-frequency cases show a small decrease in vorticity, consistent with the experiments.

At $x/D = 1.9$, the low-frequency cases have diffused the most, followed by the high-frequency cases, and the stationary case the least. This behaviour is consistent with the experiments, albeit with less magnitude differences. At $x/D = 2.5$, the order from low to high vorticity has not changed and remains comparable to the experiments.

8.3.4. Far Wake

In the far wake, the tip region can still be identified because of the positive vorticity. This diffuses further while travelling downstream. Both the simulations and the experiments show vorticity profiles with small and oscillating vorticity. The higher vorticity of the stationary case is not captured, neither are other small variations between motion cases.

8.4. Turbulence intensity

This section presents the time-averaged TI profiles at eight streamwise locations are shown in Figure 8.5. Simulation results are plotted as solid lines, experimental data as dashed lines.

8.4.1. Root Region

The TI profiles are affected by the turbulent inflow. At $x/D = 0.7$, the minimum value of TI in the nacelle jet and wake core is higher than in the base simulations, bringing it closer to the experimental results. However, the TI in the nacelle jet is still significantly underestimated due to the absence of the tower and nacelle. Above the minimum, a TI maximum appears at $z/D = 0.13$, associated with turbulence produced in the root mixing layer. Although this peak is narrower than in the experiments, its maximum is captured well. No differences between motion cases are visible in the simulations, which aligns with the minimal differences seen in the experiments.

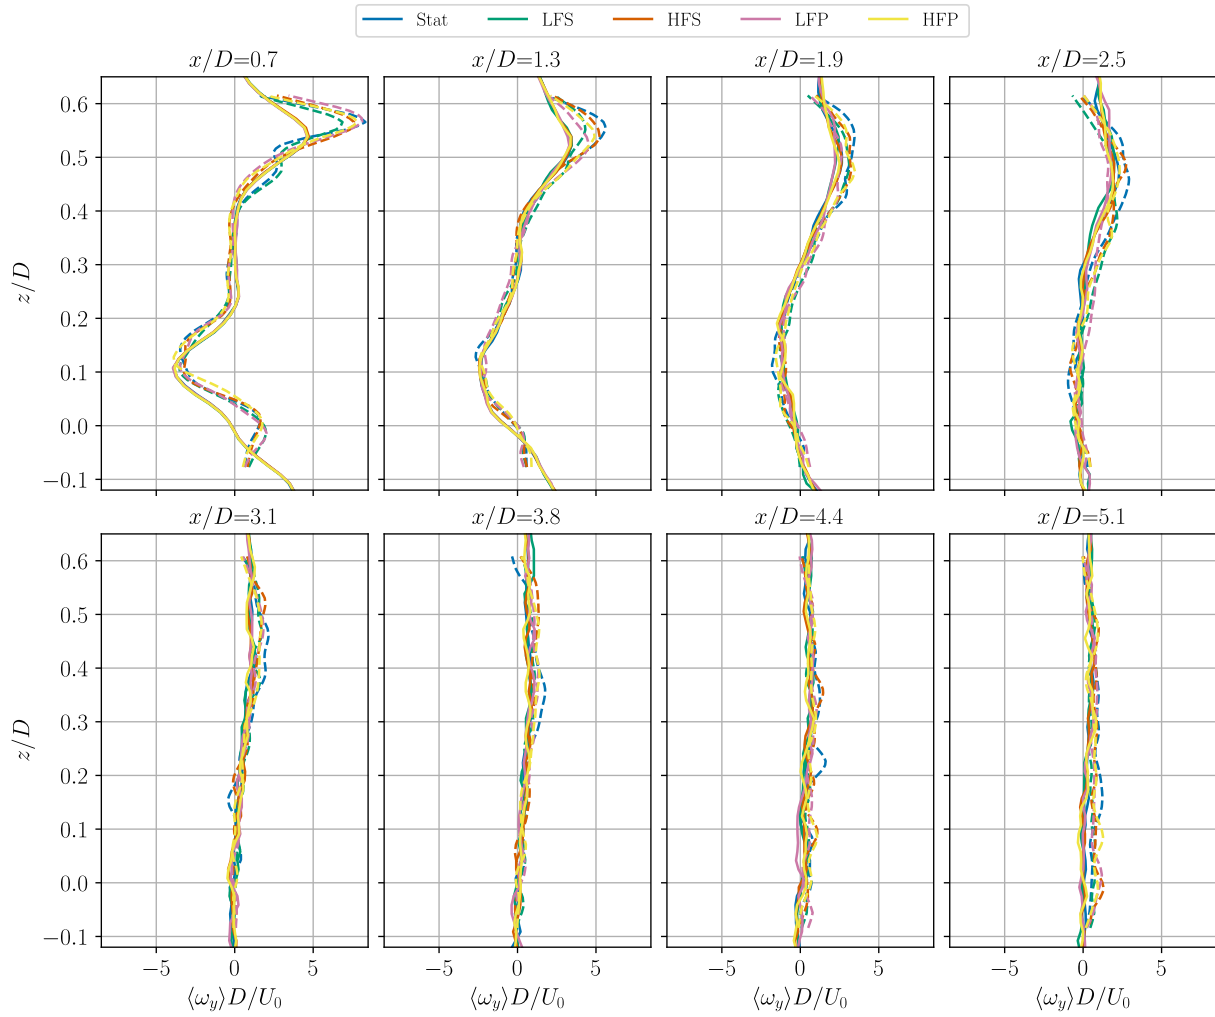


Figure 8.4: Normalised time-averaged out-of-plane vorticity along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

At $x/D = 1.3$, the TI in both the root vortex and nacelle jet is overpredicted. In the root vortex, the gradients do not match the experiments, while in the nacelle jet they do. Small differences between motion cases begin to emerge: the LFP case shows, consistent with the experiments, the lowest TI in the nacelle jet; the others are indistinguishable. In the root vortex, the differences remain negligible.

By $x/D = 1.9$, the nacelle jet exhibits a significant increase in TI, which remains overpredicted. The root mixing layer has expanded, and its TI magnitude has increased slightly. In the nacelle jet, the low-frequency cases show the lowest TI, contradicting the experiments, while the other cases are too similar to distinguish. Within the root mixing layer, the HFS case has the lowest TI and the LFS case the highest. These small differences make it difficult to identify clear consistencies with the experiments.

At $x/D = 2.5$, the TI in the root region is almost uniform at 20%, whereas the experiments show a moderate gradient with large differences between the motion cases in the TI magnitude. Between the experiments and the simulations there exists a very large difference in TI magnitude for the LFS case, followed by the LFP case and high-frequency cases, while the stationary case is predicted quite well.

8.4.2. Wake Core

At $x/D = 0.7$, the wake core shows less motion case variation than in the base configuration simulation. The stationary case has the lowest TI, followed by the high-frequency cases, and then the low-frequency cases. Experiments confirm that the stationary case is indeed the lowest and the LFP case the highest, but the ordering of the intermediate cases is different. The simulated wake core is also wider than in the

experiments due to narrower root and tip vortex TI maxima.

At $x/D = 1.3$, the simulations still show a clear and large wake core with the same order of high-to-low TI as before but with increased differences between them. These differences are consistent with the experiments. In contrast, experiments indicate that the low-frequency cases have almost no wake core remaining because of tip mixing layer growth. In a lesser extent this does also happen for the other cases. In the simulations it takes until $x/D = 3.1$ for the wake core TI minimum to disappear.

8.4.3. Tip Region

At $x/D = 0.7$, the tip region exhibits a clear TI maximum, as in the experiments, representing an improvement over the low-turbulence inflow case. The peak TI matches the stationary and LFS cases, while the other cases are underpredicted by several percent. Simulations also show differences not present in the experiments, such as the LFS case having the lowest TI. Additionally, the increased TI in the stationary and LFS cases due to the additional vortex at $z/D = 0.45$ is not modelled.

At $x/D = 1.3$, the vertical position of the maximum begins to deviate from the experiments. The simulations place the maximum at $z/D = 0.54$, whereas the experiments show a motion-dependent downward shift: low-frequency cases around $z/D = 0.42$, high-frequency cases at $z/D = 0.48$, and the stationary case at $z/D = 0.52$. The magnitude of the maximum remains underpredicted, particularly for the LFP case, while the stationary case is best represented, though slightly overpredicted. Differences between motion cases remain inconsistent with the experimental results; the low-frequency cases have the lowest turbulence and the high-frequency cases the highest in the simulations.

Toward $x/D = 2.5$, these inconsistencies persist and the vertical position mismatch grows. The ordering of the cases improves marginally. It can be seen that the stationary case has the lowest TI, followed by the high-frequency cases, then the low-frequency cases.

8.4.4. Far Wake

In the far wake, the TI profiles become approximately Gaussian. At $x/D = 3.1$, the TI is nearly constant around the wake axis in both simulations and experiments, but is underpredicted in the simulations. The motion cases show varying TIs in the wake core, without consistent agreement with the experiments. Above $z/D = 0.5$, the simulations show a decreasing TI, where differences between motion cases are reasonably well captured: LFP has the lowest TI, followed by LFS, then the stationary and high-frequency cases. It appears in this region, the gradient is also correct, however, the vertical position is not.

Toward the end of the domain, the TI increases for all cases in both simulations and experiments, but the underprediction near the wake axis persists. The vertical mismatch increases and differences between motion cases become less consistent due to small fluctuations changing their ordering along the profile.

8.5. Reynolds stress

This section presents the normalised time-averaged streamwise-vertical Reynolds stress profiles at eight streamwise locations are shown in Figure 7.4. Simulation results are plotted as solid lines, experimental data as dashed lines.

8.5.1. Root Region

At $x/D = 0.7$, the simulations reproduce the stress maximum observed in the experiments, but it is slightly overpredicted and shifted downward. The absence of nacelle and tower effects also causes the stress minimum below the wake axis to be displaced too far downward. Differences between motion cases are not visible at this location.

At $x/D = 1.3$, the simulated maximum increases further, whereas the experiments show a decrease. By $x/D = 1.9$, the maximum in the simulations has shifted upward, and differences between motion cases become more pronounced. This upward shift is not observed in the experiments, creating a vertical mismatch. The distribution of motion cases is only consistent with the experiments for the stationary,

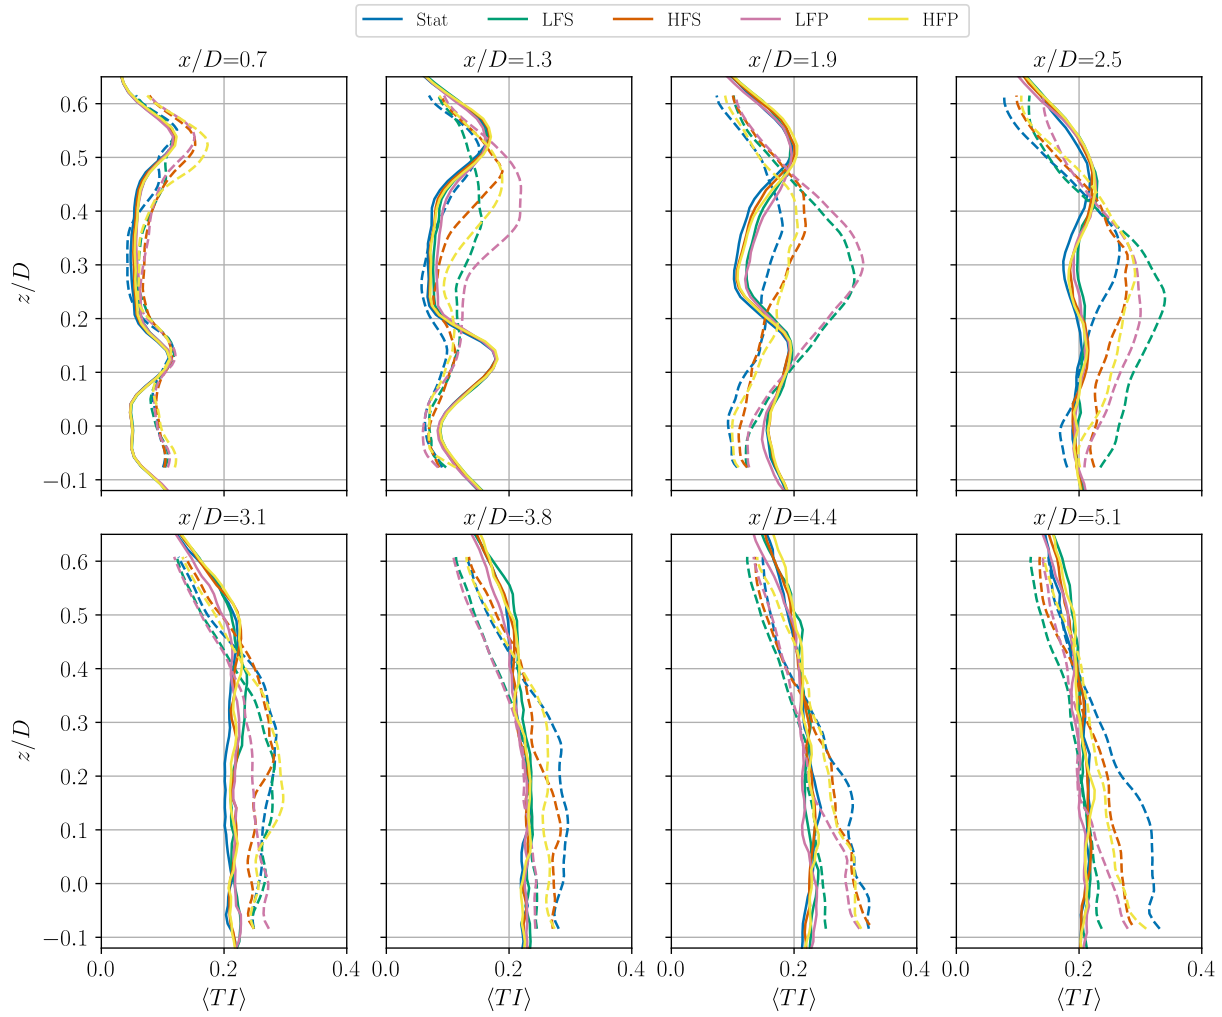


Figure 8.5: Time-averaged TI along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

high-frequency and LFP cases, with the latter having the highest stresses. The stress in the LFS case is overpredicted most.

At $x/D = 2.5$, the root stress maximum continues to shift upward in both simulations and experiments. The gradient around $z/D = 0.2$ matches well with the stationary and HFS cases in the experiments, although a vertical offset remains for the other cases. Below $z/D = 0.1$, the simulations predict steeper gradients, leading to an overestimation of the absolute stress near the wake axis. Differences between motion cases are not accurately reproduced here, but around $z/D = 0.2$ they are captured well; the low-frequency cases have the lowest stresses, followed by the stationary and high-frequency cases.

8.5.2. Wake Core

At $x/D = 0.7$, the wake core shows near-zero stresses due to the low fluctuations, extending accurately from the root vortex to the tip vortex.

By $x/D = 1.3$, the experiments show a small deviation from zero in the wake core for the low-frequency cases, which is not captured by the simulations. Nonetheless, the height of the wake core is still well represented.

At $x/D = 1.9$, the growth of the tip and root mixing layers narrows the wake core. There is a vertical shift in the core location: in the simulations it is at $z/D = 0.3$, while in the experiments it is at $z/D = 0.2$. This location is similar across all motion cases in the simulations.

By $x/D = 2.5$, the wake core has effectively vanished, becoming the interface between the remnants of the root and tip vortices. A vertical shift remains between simulations and experiments at this location.

8.5.3. Tip Region

At $x/D = 0.7$, the simulation shows a similar minimum for all motion cases, corresponding closely to the LFP and high-frequency cases. The stationary and LFS cases are overpredicted in absolute stress, and the high concentration of the stress observed in the experiments is not captured.

By $x/D = 1.3$, the stress increases. The simulation matches the LFP case almost perfectly, while the other cases are overpredicted, as the simulation predicts only minor differences between motion cases.

At $x/D = 1.9$, the stress minimum to shift downward in the experiments. This downward shift is underestimated in the simulations. The relative differences between motion cases are consistent with the experiments: higher absolute stress in the low-frequency cases, followed by the high-frequency and stationary cases. However, the differences are smaller than in the experiments, and the absolute stress is generally overpredicted except for the LFP case.

By $x/D = 2.5$, the LFP case shows a large decrease in stress in the experiments, which is not observed in the simulations. The LFS case is predicted accurately in terms of absolute stress, but its vertical position differs significantly. The high-frequency case shows both a vertical shift and an accurate prediction of absolute stress, while the stationary case is overpredicted.

8.5.4. Far Wake

At $x/D = 3.1$, the simulations show negative stress throughout, whereas the stationary, LFS, and high-frequency experimental cases remain slightly positive. The profiles appear to have a vertical mismatch with the experiments. The stresses below $z/D = 0.3$ are well predicted, but overpredicted above this point. The order of motion cases varies significantly throughout the profiles, similar to the experiments. Generally, around the wake axis the LFP case is lowest, followed by the high-frequency and stationary cases, and the LFS highest. In the top part of the wake, the stationary and high-frequency cases have the most negative stresses and thus the low-frequency cases the lowest. Here, the LFS does not comply with the experiments since in the simulations it has the most negative stresses.

By $x/D = 3.8$, the absolute stresses are overpredicted along most of the profile, especially for the LFS case. The relative order of cases along the profile do not change significantly. It appears that the high-frequency cases are close to the stationary case which is similar to the experiments, but their order is accurate everywhere.

Towards the end of the domain, the stresses start to become more uniform. The vertical mismatch increases further and differences between motion cases increase. These differences vary in consistency with the experiments.

8.6. Vertical Velocity

Comparisons between simulations and experiments reveal a significant vertical shift in the wake that develops downstream, motivating a closer investigation of vertical velocity profiles.

Figure 8.7 presents the normalised time-averaged vertical velocity profiles along the vertical direction at eight streamwise locations. Simulation results are shown as solid lines, while experimental data are shown as dashed lines.

The differences between simulation and experiment are substantial. At $x/D = 0.7$, there is no distinct minimum in the nacelle jet, likely due to the absence of nacelle modelling. Towards the top of the wake, the vertical velocity increases and approaches experimental values. Throughout the profiles there are negligible differences between motion cases.

By $x/D = 1.3$, the effect of the nacelle jet is still visible, but agreement improves. Around $z/D = 0.5$, however, the experiments and simulations show a downward vertical velocity. However, that of the experiments is higher.

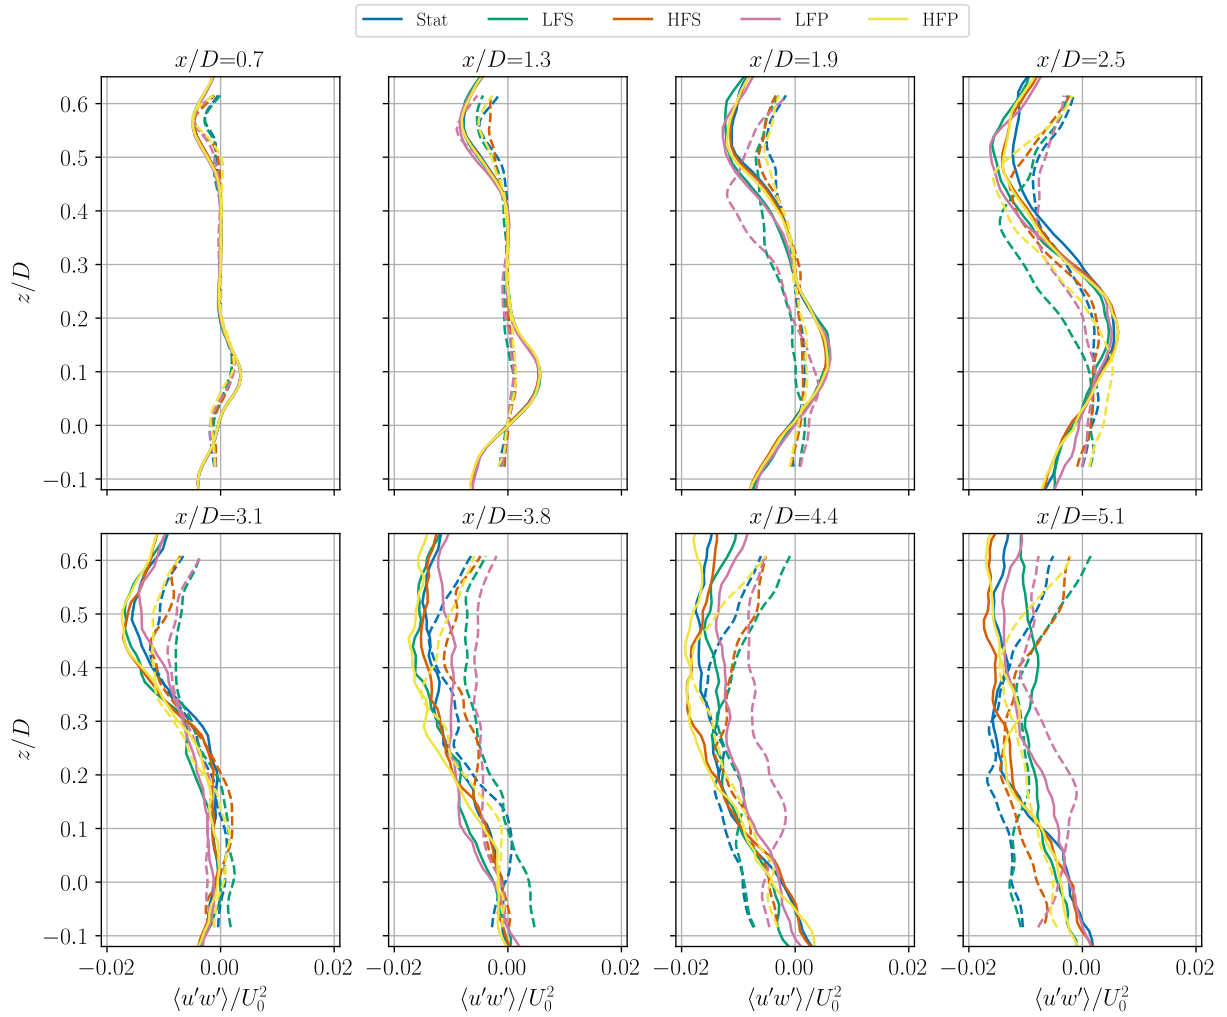


Figure 8.6: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

At $x/D = 1.9$, simulation profiles remain closely grouped, with low-frequency cases showing slightly more negative vertical velocities than the others. While the differences are smaller than in the experiments, the trend is consistent.

At $x/D = 2.5$, low-frequency cases in the experiments show a fast increase in downward velocity, which is not simulated. The other cases remain similar to the experiments. Further downstream, the underprediction of the downward velocity grows for all cases into a almost uniform velocity profile.

8.7. Linking the Metrics

To address research question 3-“How does inflow turbulence affect the wake?”-this section examines the relationships between the different metrics which were discussed in this chapter. The wake is considered in three regions: the root region, the wake core, and the far wake. A comparison is made between the base configuration and the turbulent inflow configuration, and between the simulations and experiments. Additionally, the vertical profile shift which was observed is discussed.

8.7.1. Root Region

The nacelle jet streamwise velocity remains consistently overpredicted due to the absence of the nacelle and tower in the simulations. This has not changed due to the turbulent inflow. In contrast, the root streamwise velocity gradient is captured well because the root vorticity magnitude is more accurately modelled, although its position is slightly shifted downward compared to the experiments. This could

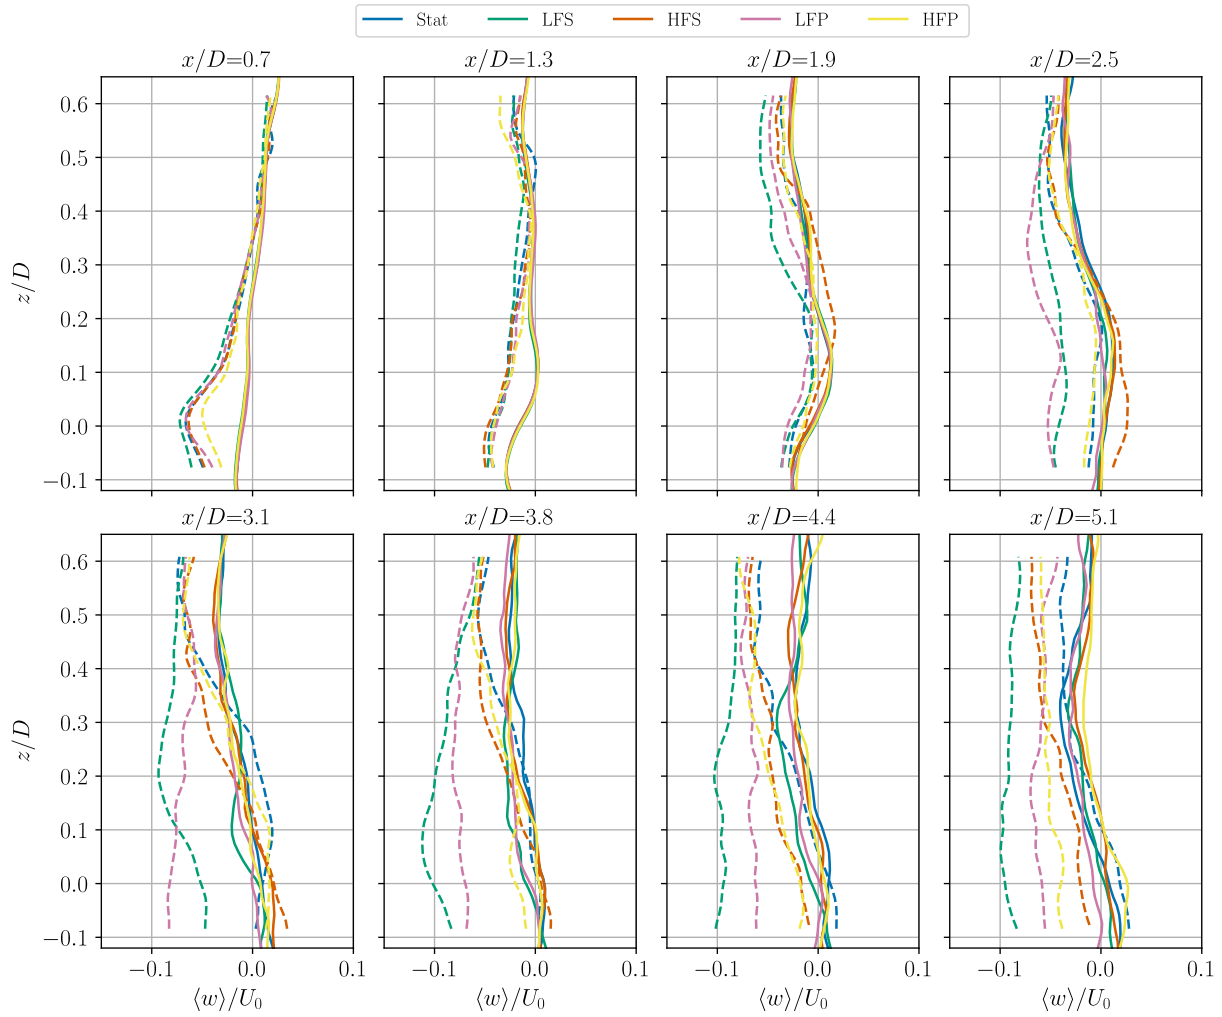


Figure 8.7: Normalised time-averaged vertical velocity along the vertical direction from the base simulation (solid line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

indicate a difference in inflow condition. As a result the velocity gradient is positioned a bit lower than in the experiments.

Beyond $x/D = 1.3$, the root region the tower and nacelle effects reduce as the vorticity diffuses. This process is faster than with the low-turbulent inflow. The cause is the increased Reynolds stresses and TI, although they are overpredicted. Without the turbulent inflow the stresses and TI were underpredicted.

8.7.2. Wake Core

At $x/D = 0.7$, the turbulent inflow causes the TI minimum to increase towards values that are consistent with the experiments. The reason is that the rotor turbulence is superimposed with the inflow turbulence. The differences between the motion cases were already accurate without turbulent inflow. It is observed that it is still the case, but with smaller differences due to the added turbulence.

The tip vortices and streamwise velocity profiles show that the wake expansion stops before $x/D = 0.7$, consistent with the experiments and an improvement on the base configuration which showed an excessive expansion.

The tip vorticity is underpredicted more than in the base simulations. The reason is that there is significantly more turbulence in this region which diffuses the vorticity quickly. This diffusion is accurate however, since the diffusion rate matches the experiments, resulting in a consistent streamwise velocity evolution. this is a major improvement on the base simulations. Vorticity differences between motion

cases remains small though. Therefore, keeping the difference in streamwise velocity small too in the near wake. Some differences are consistent with the simulations.

8.7.3. Far Wake

The initial overprediction of Reynolds stress decreases and begins to fall below the experimental values after $x/D = 1.9$. This improved agreement with the experiments enhances the comparison between vorticity and velocity. However, the TI, which is significantly higher in the wake core in the experiments, starts to dominate the wake recovery beyond $x/D = 3.1$. At this point, the streamwise velocity becomes underpredicted by some cases, having previously been overpredicted by them.

The differences between motion cases are captured with varying accuracy. At some locations, the ordering of stationary, high-frequency, and low-frequency cases is reproduced reasonably well. However, these differences are small and not consistently observed across all positions. Differences between the surge and pitch motion for the same frequency are usually predicted inaccurately.

8.7.4. Vertical Shift of the Wake

As discussed in the previous sections, a vertical shift is observed throughout the wake, particularly beyond $x/D = 1.9$. The vertical velocity profiles reveal a mismatch between the simulations and experiments after this point. Towards the downstream end of the measurement domain, the simulations also show a downward velocity, although its magnitude is much smaller than in the experiments. The origin of this additional downward velocity remains unclear, as noted in Section 3.4.1.

An attempt was made to adjust the simulation profiles for this shift based on the downward velocity throughout the wake, as described in Appendix D. However, this method was not applied to the results in the next chapter, as the underlying physics can only be hypothesised. Possible explanations include the effect of non-buoyant helium–soap particles, the presence of a low-pressure region below the wake drawing it downwards, or a deficiency in the simulation itself.

FALM, Nacelle and Resolution

To further improve the simulations this chapter presents an analysis of various simulation with an adjusted configuration. First, the effect of the FALM will be analysed in Section 9.1 to answer research question 4. Secondly, the drag of the nacelle is added to the simulation in Section 9.2 to answer research question 5.

9.1. FALM

In the previous chapter, it was shown that the tip vorticity remains underpredicted by the simulations, as was also observed in the base simulations. In Section 7.5, it was suggested that the FALM could improve vorticity modelling, since this method corrects the downwash and, consequently, the load distribution for a suboptimal kernel size. A higher load gradient at the tip would strengthen the tip vortex, thereby reducing the vorticity underestimation.

However, using this method instead of the conventional ALM did not produce any differences at the current mesh resolution. This is evident from the errors in the normalised time-averaged out-of-plane vorticity between the FALM simulations and the experiments (solid lines) in Figure 9.1. For comparison, the figure also includes the vorticity error profiles from the turbulent inflow simulation using the ALM (dotted lines). The other metrics have been included in Section A.4.




Figure 9.1: Normalised time-averaged out-of-plane vorticity difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All differences are along the vertical direction of the domain for all motion cases. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.

9.2. Including Nacelle

It was observed that the nacelle and tower have a pronounced effect on reducing the flow speed below $z/D = 0.05$ until $x/D \approx 1.3$. The lower part of the root vorticity is displaced upwards, while the TI and Reynolds stress increase. Further into the near wake, the experiments show differences between motion cases in the root region that are absent in the previously presented simulations. This may be partly due to the omission of the nacelle and tower in the model. To improve the simulation, it is therefore necessary to include the nacelle.

9.2.1. Velocity Sampling Point

As explained in Section 4.3.3, the nacelle is modelled using a single actuator drag point. The drag coefficient and frontal area are 1.05 and 26 cm², respectively. The position of the velocity sampling point must be placed upwind of the rotor, at a distance that ensures the nacelle forcing is not applied directly at the cells where the force is applied, but where the flow is decelerated by the rotor. The required upwind distance is calculated by determining where the kernel weight drops to 1% of its maximum, as shown in Equation 9.1, where w is set to 1% and the kernel size ϵ is the same as for the blades—three times the cell size of 0.017 D . This yields an upwind distance of 0.23 D .

$$w_f = e^{-\frac{r}{\epsilon}} \Rightarrow r = -\epsilon \ln w_f \quad (9.1)$$

9.2.2. Nacelle Effect

The improvements from adding the nacelle are small for both the out-of-plane vorticity and the streamwise velocity, which are directly related via the induced velocity. The TI and streamwise-vertical Reynolds stresses show no significant changes, as the nacelle drag point does not generate turbulence in the absence of simulated surface shear stresses. These metrics are therefore not discussed here, but are included in Section A.5.

In Figure 9.2, the normalised time-averaged streamwise velocity profiles from simulations including the nacelle are shown at four streamwise locations behind the rotor plane. The far wake ($x/D > 2.5$) is excluded since the nacelle's influence has fully dissipated by this point. The nacelle effect is identified by comparing these profiles with those from Figure 8.3, where no nacelle was included.

At $x/D = 0.7$, adding the nacelle primarily affects the flow below $z/D = 0.05$, where it reduces the velocity overprediction by introducing drag. However, a substantial portion of the overprediction remains because the tower is not modelled and the nacelle is permeable, allowing flow to pass through it.

Further downstream, the nacelle continues to cause a small velocity reduction, but its influence gradually diminishes. By $x/D = 2.5$, no measurable difference remains compared to the nacelle-free case.

As can be observed, the differences between the motion cases do not increase in the nacelle jet. This is because the differences are mostly related to the differences in shear stresses and the resulting turbulent diffusion.

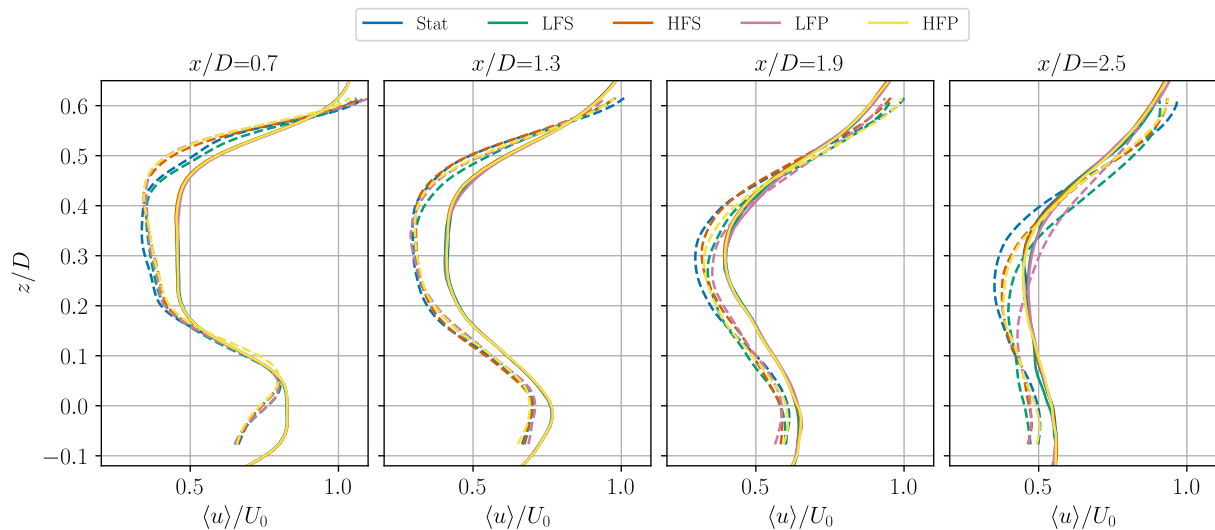


Figure 9.2: Normalised time-averaged streamwise velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to . 0.7, 1.3, 1.9 and 2.5.

The reduced streamwise velocity in the nacelle jet is linked to the decrease in root out-of-plane vorticity shown in Figure 9.3. Here, the overprediction of the root vorticity magnitude is reduced, and the gradient below $z/D = 0.1$ matches the experiments more closely. This yields a more accurate induced velocity

and an improved near-wake representation.

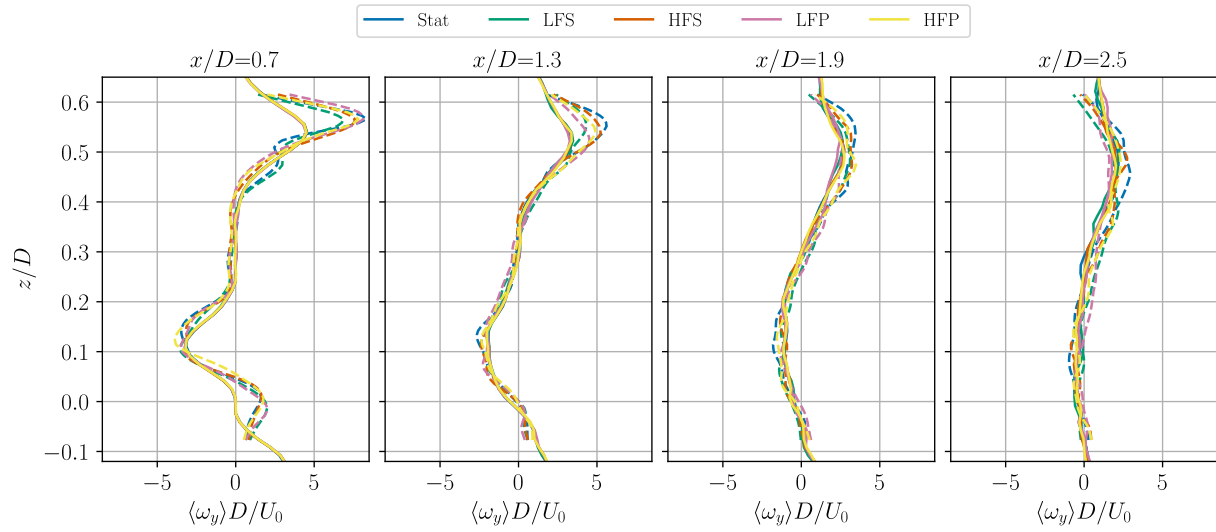


Figure 9.3: Normalised time-averaged out-of-plane velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to . 0.7, 1.3, 1.9 and 2.5.

A similar improvement would be achieved by adding the tower to the simulations. However, difficulties were encountered and time constraints resulted in not completing these simulations. What has been done is presented in Appendix C.

High Resolution

In Section 5.5, resolution convergence was assessed using the streamwise velocity profile without turbulent inflow and without nacelle. However, the observed underprediction of out-of-plane tip vorticity, which could be the cause of the overprediction of streamwise velocity in the wake core, indicates that the mesh was under-resolved. This chapter presents results from simulations with turbulent inflow and a nacelle, using a higher mesh resolution of $\Delta x/D = 0.012$ to improve vorticity modelling.

A higher resolution was also expected to enhance the modelling of differences between motion cases, but the results show this is not the case. The following analysis therefore focuses solely on the general effect of resolution on the profiles, excluding any assessment of motion case differences.

In Section 10.1, the inflow conditions are compared with those from the coarser simulation to verify their similarity. The vertical velocity profiles are discussed in Section 10.2 to assess the vertical displacement of the wake. The simulation results are then evaluated against the established comparison metrics in Section 10.3, Section 10.4 and Section 10.5. Reynolds stress is omitted from the main text and presented in Section A.6, as its observations are similar to those of the TI analysis. Finally, the interaction between metrics is examined in Section 10.6.

10.1. Inflow

The high-resolution simulations required a new inflow field. Its main profiles are shown in Figure 10.1 alongside those from the coarse simulations. Minor differences arise because turbulence synthesis is inherently stochastic, but using similar settings and allowing sufficient development time produces an inflow closely matching the original. The most notable difference is an offset in the spectra, caused by the finer mesh resolving more energy than the coarse mesh.

10.2. Vertical Velocity

Figure 10.2 presents the normalised time-averaged vertical velocity profiles from the simulations (solid lines) and experiments (dashed lines) at eight streamwise locations behind the rotor plane. These results are compared with the coarser simulation profiles in Figure 8.7.

At $x/D = 0.7$, the high-resolution simulations exhibit a higher vertical velocity than the coarse case, attributable to the less negative inflow velocity in the upper part of the wake. However, this does not explain the behaviour below $z/D = 0.1$, where the inflow velocity is more negative, and the cause remains unclear.

Further downstream, the downward velocity increases in both resolutions but does so more rapidly in the high-resolution case. Consequently, after $x/D = 3.1$ the high-resolution simulations produce a greater downward velocity, reducing the mismatch with the experimental data compared with the coarse simulation.

10.3. Streamwise Velocity

The impact of higher resolution on the normalised time-averaged streamwise velocity is shown in Figure 10.3, which presents profiles at eight streamwise locations behind the rotor plane. Solid lines represent simulations, and dashed lines represent experiments.

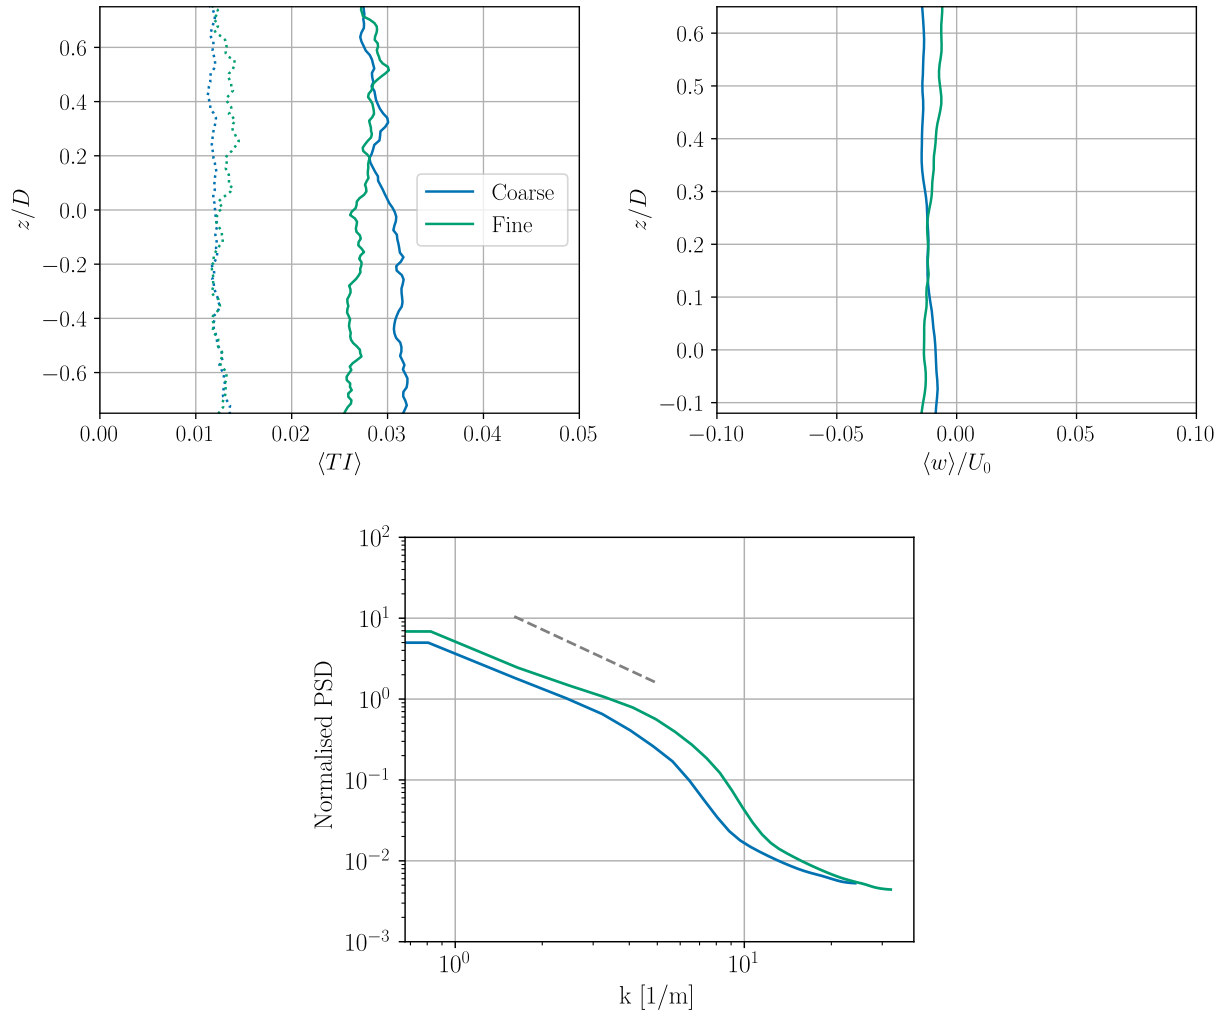


Figure 10.1: High-resolution precursor inflow conditions compared to the inflow conditions from the low-resolution precursor. Top-left is the time-averaged TI (solid lines) and its standard deviation (dotted lines). Top-right is the normalised time-averaged vertical velocity. Bottom are the spatial turbulence spectra.

10.3.1. Root Region

At $x/D = 0.7$, the higher resolution results in faster flow above and below the nacelle. This is due to a reduced sphere of influence, as the absolute kernel size decreases with cell size. Numerically increasing the nacelle's frontal area could compensate for this effect.

Further downstream, the overprediction of streamwise velocity persists until nacelle and tower effects dissipate at $x/D = 3.1$. Beyond this point, high- and low-resolution simulations converge around the wake axis.

10.3.2. Wake Core

At $x/D = 0.7$, the wake core velocity remains overpredicted, as before. However, the velocity gradient in the tip region is positioned higher, improving height of the core. The velocity fluctuation at $z/D = 0.25$ is more pronounced due to the smaller kernel size (see Section 5.4) but disappears by $x/D = 1.3$.

By $x/D = 2.5$, the wake core is positioned even further upward compared to the low-resolution case, consistent with the increased mismatch in vertical velocity discussed in Section 10.2.

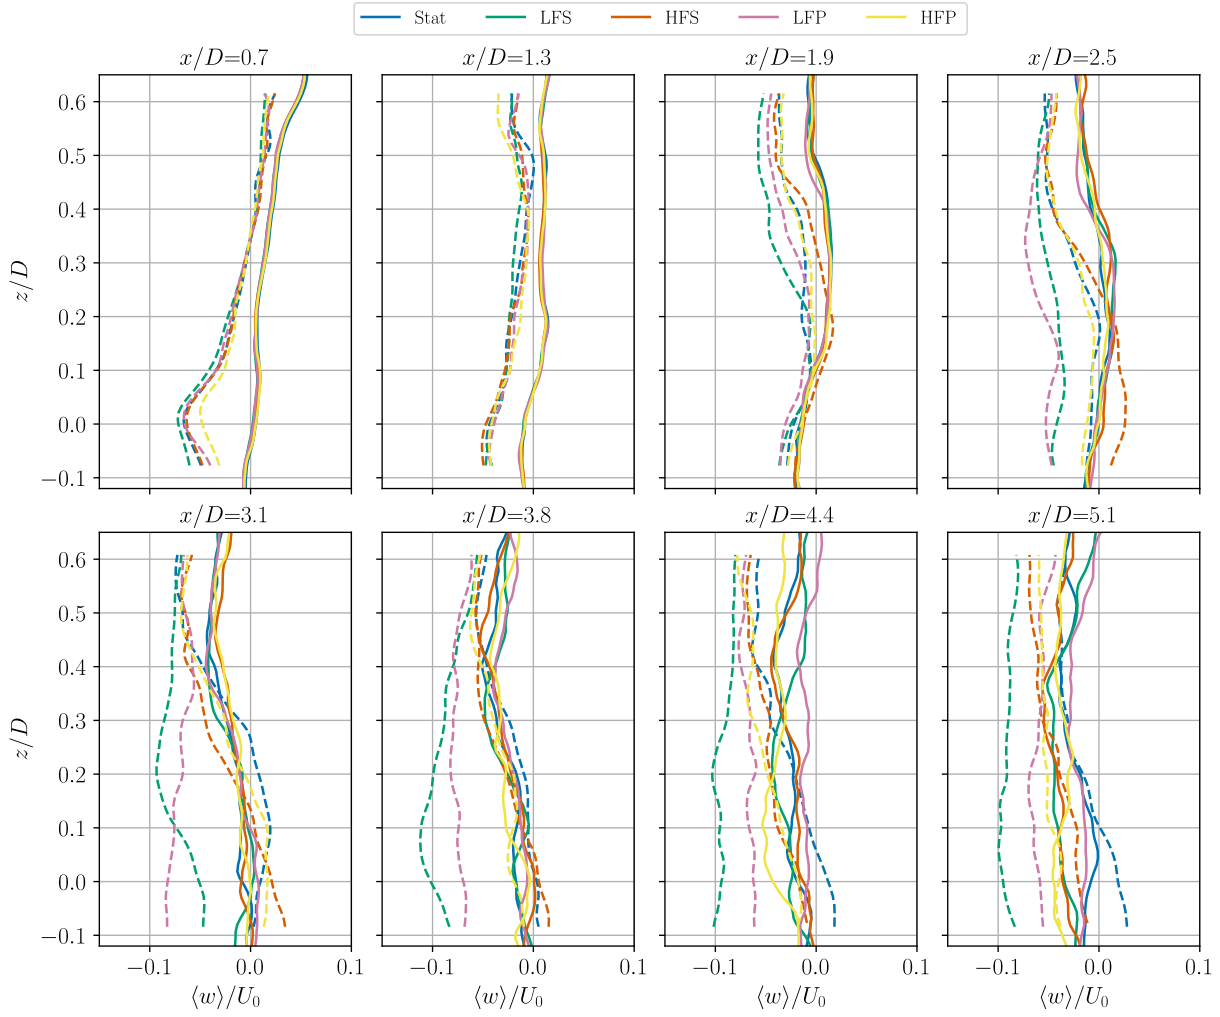


Figure 10.2: Normalised time-averaged vertical velocity along the vertical direction from the base simulation (solid line) for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. Additionally, the profiles at the same locations in the experiments have been included (dashed line).

10.3.3. Tip Region

At $x/D = 0.7$, the location of the tip velocity gradient is more accurately captured, and its magnitude appears improved.

Downstream, the gradient shifts upward, creating an offset from the experimental results due to the vertical velocity mismatch, peaking at $x/D = 2.5$.

10.3.4. Far Wake

In the far wake, the higher resolution introduces small differences between motion cases, linked to variations in the TI field. However, their relative behaviour remains unchanged. The vertical velocity mismatch with experiments is slightly reduced here compared with the low-resolution case, improving agreement with measurements.

10.4. Vorticity

The impact of higher resolution on the normalised time-averaged out-of-plane vorticity is shown in Figure 10.4, which presents profiles at four streamwise locations behind the rotor plane. Solid lines indicate simulations, and dashed lines indicate experiments. The far wake is omitted, as the higher resolution produces minimal changes there aside from the vertical shift of profiles relative to experiments. This behaviour is consistent across other metrics.

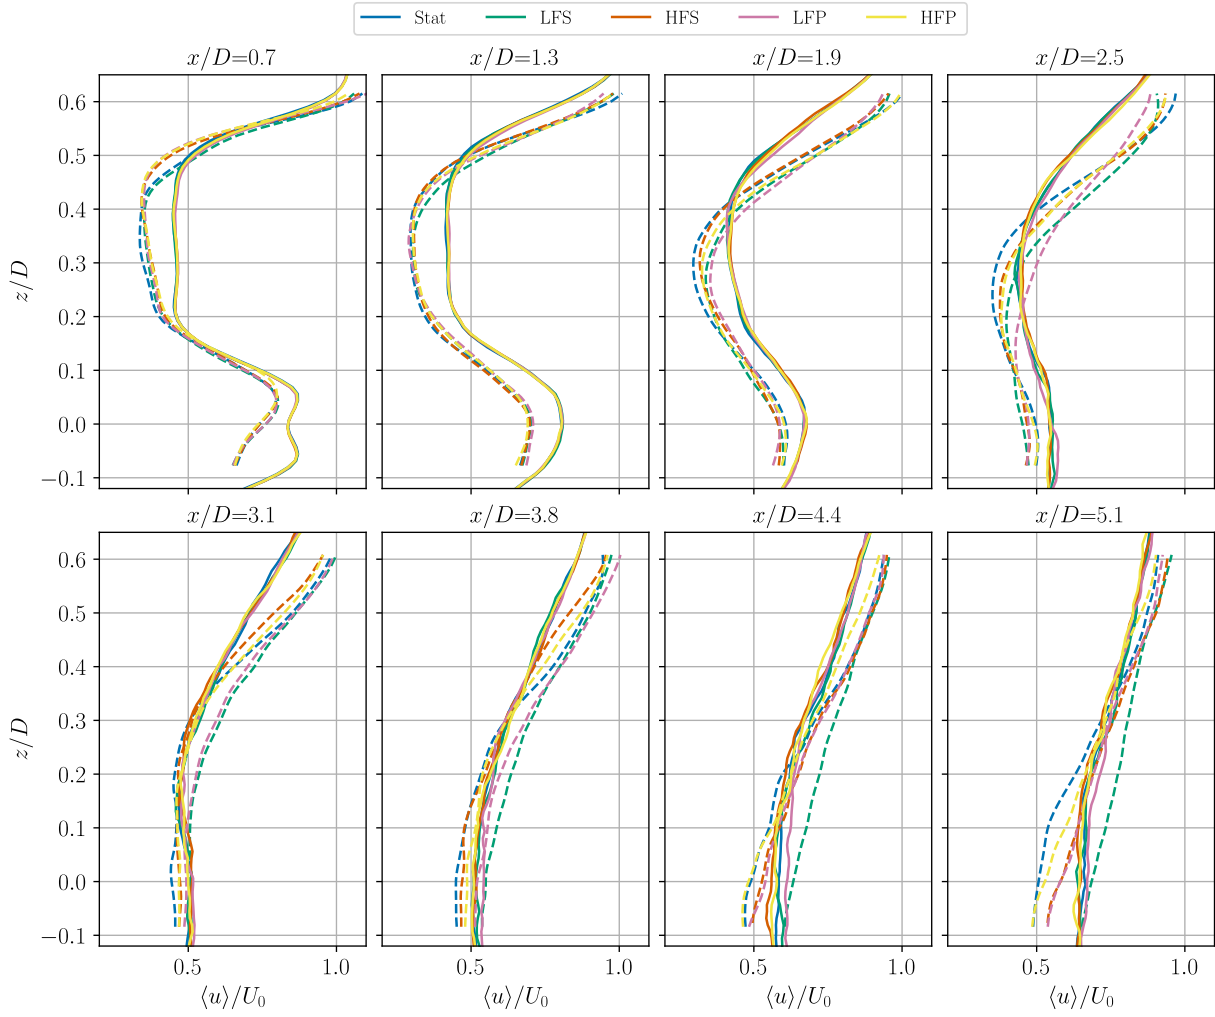


Figure 10.3: Normalised time-averaged streamwise velocity along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

10.4.1. Root Region

At $x/D = 0.7$, the vorticity of the main root vortex at $z/D = \pm 0.12$ is higher than in the low-resolution simulations, now overpredicting experimental values except in the HFP case. Additional vorticity extremes appear at the top and bottom of the nacelle, improving the match in the upper root vortex gradient compared with the low-resolution case. However, the lower extreme at $z/D = -0.04$ increases the discrepancy with experiments; a tower would likely suppress this minimum through turbulent diffusion. By $x/D = 1.3$, differences from the low-resolution case are negligible.

10.4.2. Tip Region

At $x/D = 0.7$, the tip vortex strength is markedly higher, with its position more accurately captured. By $x/D = 1.3$, the tip vortex is shifted upward in the high-resolution simulations due to the overly positive vertical velocity noted earlier, increasing the vertical mismatch with experiments. However, the peak vorticity magnitude is better reproduced than in the low-resolution case. By $x/D = 1.9$, the only remaining difference is this vertical offset.

10.5. Turbulence Intensity

The impact of higher resolution on the normalised time-averaged TI is shown in Figure 10.5, which presents profiles at four streamwise locations in the near wake. Solid lines represent simulations, and dashed lines represent experiments.

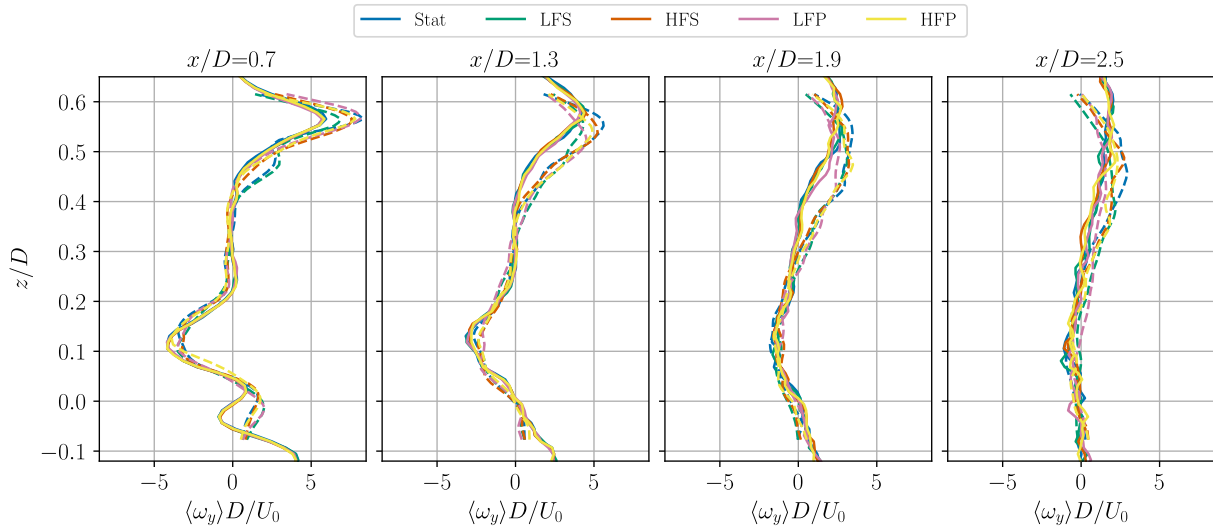


Figure 10.4: Normalised time-averaged out-of-plane vorticity along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.

10.5.1. Root Region

At $x/D = 0.7$, TI within the nacelle jet increases slightly, while the maximum at $z/D = 0.12$ decreases marginally. A similar pattern appears at $x/D = 1.3$, where higher resolution reduces TI within the root vortex, improving agreement with experiments, though some overprediction remains.

By $x/D = 1.9$, overprediction persists for the stationary and high-frequency cases but is reduced compared with the low-resolution simulations. The TI increase observed in the low-frequency experimental cases is still not captured. Around $x/D = 2.5$, TI in the root region remains higher than in the low-resolution case, but this improvement diminishes quickly downstream.

10.5.2. Wake Core

In the near wake, turbulence fills the wake core more rapidly in the high-resolution simulations. At $x/D = 0.7$, differences from the low-resolution case are negligible. By $x/D = 1.9$, TI in the core increases due to the higher resolution, with a similar increase at $x/D = 2.5$. In the far wake, the difference becomes negligible again.

10.5.3. Tip Region

As in the root region, TI in the tip region decreases at $x/D = 0.7$ with higher resolution, and this reduction persists until $x/D = 2.5$. Further downstream, the high- and low-resolution results converge.

10.6. Linking the Metrics

The aim of increasing the resolution was to assess whether it would reduce the underprediction of out-of-plane vorticity in the tip region, and whether this increase in vorticity would mitigate the overprediction of streamwise velocity in the wake core. The results show that, partly due to reduced numerical diffusion, both root and tip vorticity magnitudes increase, and the tip vorticity is located more accurately in the vertical direction. The higher resolution also increases the blade load gradient at the root and tip, as shown in Figure 10.6. Here, the solid lines represent the projected thrust force per unit volume along the rotor radius for the high-resolution simulations, while the dashed lines show the low-resolution results. The blade load gradient is directly linked to bound circulation variation, and thus to tip vorticity magnitude, as discussed in Section 2.1.3. Therefore, the higher gradient also contributes to the vorticity increase.

From the streamwise velocity profiles, this improvement is seen only in the magnitude and location of the tip velocity gradient. The overprediction of wake core streamwise velocity remains unchanged, suggesting the cause may be linked to the lack of blade geometry and impermeability modelling, neglecting

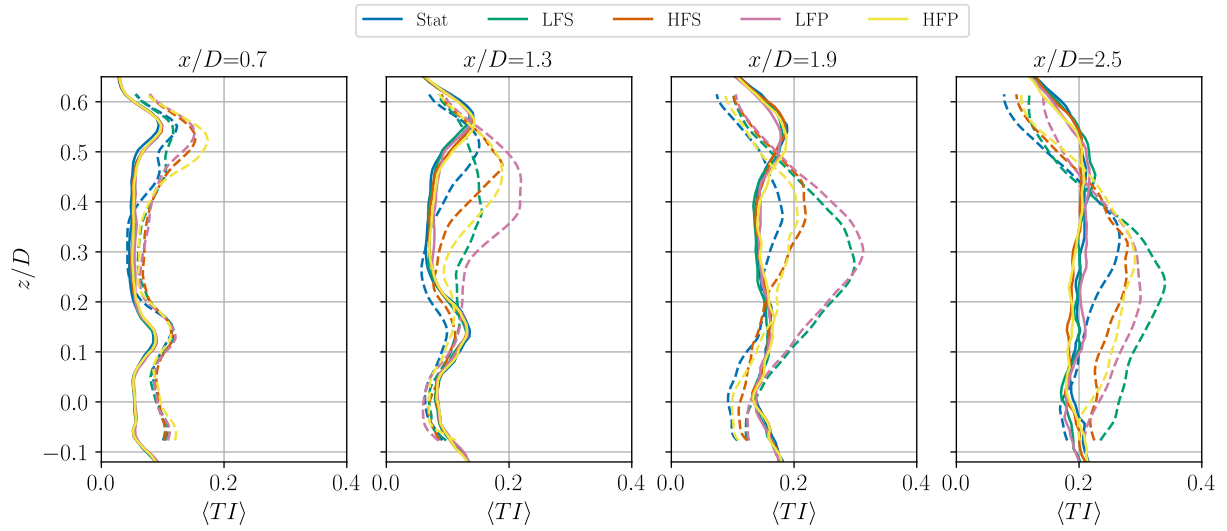


Figure 10.5: Time-averaged TI along the vertical direction from the simulations with a high resolution for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.




Figure 10.6: Time-averaged thrust per unit volume along the actuator line for the high-resolution simulations (solid lines) and low-resolution simulations (dashed lines).

three-dimensional aerodynamic effects (see Section 2.1.2) Confirming this would require higher-fidelity simulations with resolved blades.

In the near wake, vertical velocity is more strongly overpredicted in the high-resolution simulations due to differences in the inflow, resulting in an increased upward shift in the profiles of other metrics. The correction method in Appendix D could help to correct of the vertical mismatch.

The reduced absolute kernel size also decreases the sphere of influence of the nacelle, which could be counteracted by increasing its numerical frontal area.

The TI is slightly lower in the root and tip regions, but the differences are small, meaning turbulent diffusion of vorticity, and thus wake recovery, is only minimally affected. In the far wake, there are effectively no changes due to the increased resolution. The higher resolution does not improve the modelling of the sudden TI increase observed for the low-frequency cases.

Part IV

Conclusion

Conclusions and Future Scope

This chapter presents the conclusions in Section 11.1. After that, recommendations are made for future work in the research area in Section 11.2

11.1. Conclusions

FOWTs offer a promising solution for accessing deep-water wind resources. While research in this field is extensive, important questions remain unanswered. A major challenge lies in understanding and modelling the wake aerodynamics of these machines, which differ from their fixed-bottom counterparts due to the added complexity of floater motion.

Literature suggests that wake behaviour can depend on the motion DOF, frequency, and amplitude. However, the specific effects of different motions on wake evolution remain unclear. This understanding is crucial for wind farm design, as farm layouts rely on accurate predictions of wake length and turbulence intensity to model the energy capture potential and the life span of the turbines.

At Delft University of Technology, experimental campaigns with a scale-model wind turbine using a HIL approach aim to investigate the influence of motion on turbine performance and wake development. To complement this research, an accurate numerical model is required. The focus of this work is therefore to determine the optimal configuration of the AspFAST model, with the aim:

”To numerically replicate the wake aerodynamics of a scale wind turbine model with prescribed sinusoidal surge and pitch motion, using a Large Eddy Simulation (LES) Actuator Line Method (ALM) in AspFAST.”

This aim is reached by answering a set of research questions. These questions will be answered below.

1. Which configuration options in AspFAST affect the wake development, and how?

In Chapter 5, a parameter study showed that the numerical results are affected much by the domain size. If the turbine is too close to the upwind boundary, the flow cannot naturally decelerate towards the turbine. Also, the downwind distance should be large enough to capture the wake of the turbine without any influence of the downwind boundary. The domain height and lateral extent of the domain must be large enough to prevent blockage, which would otherwise increase the streamwise velocity in the wake.

Resolution was important for resolving streamwise velocity gradients in the root and tip regions. Convergence of these gradients was observed, but Chapter 10 showed that higher resolution improved tip-vorticity modelling.

Because the domain is periodic, the turbine wake can re-enter the inlet. This is mitigated by applying nudging at the boundaries to restore a uniform inflow. The nudging strength was found to be a critical parameter for ensuring wake accuracy.

Increasing the number of actuator points improved the representation of blade load distribution, with results converging at around 50 points, consistent with literature.

The kernel size affected velocity gradients: larger values smoothed gradients excessively, while too small a value caused spurious fluctuations at low resolution. A resolution-relative kernel size of 3 was chosen as a compromise; for high-resolution meshes, a value of 4 cell widths may be more suitable.

2. What are the characteristics of the experimentally measured wake, and how do they differ between motion cases?

The experiments, analysed in Chapter 6, examined the normalised time-averaged streamwise velocity,

out-of-plane vorticity, TI, and streamwise–vertical Reynolds stresses. The results reveal the typical wake features: nacelle jet, root vortices, wake core, tip vortices, and far wake. The nacelle jet was linked to the induced velocity of the root vortices, while the combined influence of root and tip vortices formed the wake core. Along the streamwise direction, strong velocity gradients generated significant shear stresses, driving vorticity diffusion and wake recovery.

It was observed that motion enhanced wake recovery compared with the stationary turbine. Low-frequency motion cases recovered faster than high-frequency ones. Between low-frequency surge and pitch motions, clear differences were present throughout the wake, though neither consistently recovered faster. Differences between the high-frequency cases were small. The faster wake recovery in moving cases was attributed to increased Reynolds stresses, which raised the TI, particularly after tip vortex breakdown at approximately $x/D = 1.6$. In the far wake, the stationary case exhibited higher Reynolds stresses than the motion cases, potentially indicating faster recovery further downstream. However, this could not be confirmed due to the limited measurement range.

3. How does inflow turbulence affect the wake?

The influence of inflow turbulence was assessed by first simulating the wake with very low TI in Chapter 7. This was followed by the introduction of more realistic turbulence in Chapter 8.

With low-turbulence inflow, the wake remained overly stable. Root and tip vortices persisted without breakdown within the measurement domain due to a significant underprediction of stresses and TI. The effects of motion were nonetheless pronounced, and in some cases aligned with experimental observations. The most unstable cases were those with low-frequency motion, particularly the LFP case.

Introducing an inflow turbulence intensity of approximately 3.6% greatly improved wake recovery by enhancing turbulence production. However, this did not improve the overprediction of wake core velocity. Furthermore, turbulent inflow reduced the differences between motion cases in the simulations, although some variations remained consistent with the experiments.

4. How does the filtered ALM influence the wake?

The results of the simulations with the FALM were compared to simulations with the conventional ALM in Chapter 9. Its effect was found to be negligible for the resolution that was used and can thus be excluded from future simulations.

5. How does the nacelle drag affect the wake?

In Chapter 9, the nacelle was modelled as a drag disk to assess its influence on the flow. This addition reduced the streamwise velocity around the wake axis, primarily due to a decrease in root vorticity magnitude. No turbulence was generated by the nacelle, as it is not represented as a non-permeable body capable of imposing shear stresses on the flow.

Simulations at higher resolution, presented in Chapter 10, show the nacelle becoming more compact because the absolute kernel size used to distribute the nacelle drag decreases with cell size. Slightly increasing the nacelle's frontal area could compensate for this effect.

A similar improvement may be achieved by including the tower, for which preliminary steps are described in Appendix C.

11.2. Future Scope

Although the simulated wakes of the motion cases are quite consistent with the experiments, some actions can be explored to expand the knowledge gained from this work:

- Perform a blade-resolved simulation to include three-dimensional flow effects. This could possibly help to reduce the streamwise velocity overprediction. This might also decrease the underprediction of the differences between motion cases.
- Use the Verstappen SGS model with the 5th order advection scheme to reduce spurious fluctuations in the wake (see Appendix B).
- Expand the mesh generation in GRASP by allowing for local refinement. The uniform mesh makes

the simulations very expensive. To keep the benefits of parallelisation on a GPU, the multi-block octree refinement structure could be employed as is used in many Lattice Boltzmann codes, like WaLBerla¹.

- Solve the problem presented in Appendix C and include the tower to improve the velocity and vorticity in the root region of the near wake.
- Use the non-periodic boundary conditions module in GRASP to set more precise inflow conditions.
- Perform a spectral analysis of the wake of the turbine for all motion cases and compare its spectral content with that of the experiments.
- Perform an uncertainty analysis of the experiment flow fields to get a better idea of the quality of the measured flow since that assessment is completely left out in this work. The quality of the experiments is left completely out of the picture in this study to simplify the analyses.

¹<https://www.walberla.net/>

Bibliography

- [1] IRENA, “Floating offshore wind outlook,” Report, International Renewable Energy Agency, 2024. URL <https://www.irena.org/Publications/2024/Jul/Floating-offshore-wind-outlook>.
- [2] Sebastian, T., and Lackner, M. A., “Characterization of the unsteady aerodynamics of offshore floating wind turbines,” *Wind Energy*, Vol. 16, No. 3, 2013, pp. 339–352.
- [3] Manwell, J. F., McGowan, J. G., and Rogers, A. L., *Wind energy explained: theory, design and application*, John Wiley & Sons, 2010.
- [4] Leishman, G. J., *Principles of helicopter aerodynamics with CD extra*, Cambridge University Press, 2006.
- [5] Vahidi, D., and Porté-Agel, F., “A physics-based model for wind turbine wake expansion in the atmospheric boundary layer,” *Journal of Fluid Mechanics*, Vol. 943, 2022, p. A49.
- [6] Vidal, C. M., “Wake Aerodynamics of Floating Offshore Wind Turbines,” Master’s thesis, Delft University of Technology, 2025.
- [7] Taschner, E., Folkersma, M., Martínez-Tossas, L. A., Verzijlbergh, R., and van Wingerden, J., “A new coupling of a GPU-resident large-eddy simulation code with a multiphysics wind turbine simulation tool,” *Wind Energy*, Vol. 27, No. 11, 2024, pp. 1152–1172.
- [8] Heronemus, W. E., “Pollution-free energy from offshore winds,” *8th Annual Conference and Exposition*, Marine Technology Society, 1972.
- [9] Sebastian, T., and Lackner, M., “Offshore floating wind turbines-an aerodynamic perspective,” *49th AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition*, 2011, p. 720.
- [10] Sebastian, T., and Lackner, M., “Analysis of the Induction and Wake Evolution of an Offshore Floating Wind Turbine,” *Energies*, Vol. 5, No. 4, 2012, pp. 968–1000. URL <https://www.mdpi.com/1996-1073/5/4/968>.
- [11] Bak, C., Zahle, F., Bitsche, R., Kim, T., Yde, A., Henriksen, L. C., Hansen, M. H., Blasques, J. P. A. A., Gaunaa, M., and Natarajan, A., “The DTU 10-MW reference wind turbine,” *Danish wind power research*, 2013.
- [12] Hendriks, S., “Analysis of the Wake of Tandem Floating Wind Turbines,” Master’s thesis, Delft University of Technology, 2024.
- [13] Taruffi, F., Miranda Novais, F., and Viré, A. C., “An experimental study on the aerodynamic loads of a floating offshore wind turbine under imposed motions,” *Wind Energy Science*, 2024. <https://doi.org/10.5194/wes-9-343-2024>, URL <http://resolver.tudelft.nl/uuid:ee94c940-6ee0-49bf-8ed4-14c76081484b>.
- [14] Taruffi, F., Combette, R., and Viré, A., “Experimental and CFD analysis of a floating offshore wind turbine under imposed motions,” *Journal of Physics: Conference Series*, Vol. 2767, No. 6, 2024. <https://doi.org/10.1088/1742-6596/2767/6/062010>, URL <http://dx.doi.org/10.1088/1742-6596/2767/6/062010>.
- [15] Anderson, J., *EBOOK: Fundamentals of Aerodynamics (SI units)*, McGraw hill, 2011.

- [16] Du, Z., and Selig, M., *A 3-D stall-delay model for horizontal axis wind turbine performance prediction*, American Institute of Aeronautics and Astronautics, 1998. <https://doi.org/10.2514/6.1998-21>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.1998-21>.
- [17] Viterna, L. A., and Janetzke, D. C., “Theoretical and experimental power from large horizontal-axis wind turbines,” Report, NASA Lewis Research Center, Cleveland, OH (United States), 1982. <https://doi.org/10.2172/6763041>, URL <https://www.osti.gov/biblio/6763041https://www.osti.gov/servlets/purl/6763041>.
- [18] Snel, H., and Schepers, J. G., “Joint investigation of dynamic inflow effects and implementation of an engineering method,” Report, Netherlands Energy Research Foundation (ECN), Petten (Netherlands), 1995.
- [19] Crespo, A., Hernandez, J., and Frandsen, S., “Survey of modelling methods for wind turbine wakes and wind farms,” *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, Vol. 2, No. 1, 1999, pp. 1–24.
- [20] Gómez-Elvira, R., Crespo, A., Migoya, E., Manuel, F., and Hernández, J., “Anisotropy of turbulence in wind turbine wakes,” *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 93, No. 10, 2005, pp. 797–814. <https://doi.org/10.1016/j.jweia.2005.08.001>, URL <https://www.sciencedirect.com/science/article/pii/S016761050500084X>.
- [21] Lignarolo, L. E. M., Ragni, D., Krishnaswami, C., Chen, Q., Simão Ferreira, C. J., and van Bussel, G. J. W., “Experimental analysis of the wake of a horizontal-axis wind-turbine model,” *Renewable Energy*, Vol. 70, 2014, pp. 31–46. <https://doi.org/10.1016/j.renene.2014.01.020>, URL <https://www.sciencedirect.com/science/article/pii/S0960148114000494>.
- [22] Widnall, S. E., “The stability of a helical vortex filament,” *Journal of Fluid Mechanics*, Vol. 54, No. 4, 1972, pp. 641–663. <https://doi.org/10.1017/S0022112072000928>, URL <https://www.cambridge.org/core/product/434CE1AAA6E69F30D27381A14528AF5C>.
- [23] Ivanell, S., Mikkelsen, R., Sørensen, J. N., and Henningson, D., “Stability analysis of the tip vortices of a wind turbine,” *Wind Energy*, Vol. 13, No. 8, 2010, pp. 705–715. <https://doi.org/10.1002/we.391>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.391>.
- [24] Lignarolo, L. E. M., Ragni, D., Scarano, F., Simão Ferreira, C. J., and van Bussel, G. J. W., “Tip-vortex instability and turbulent mixing in wind-turbine wakes,” *Journal of Fluid Mechanics*, Vol. 781, 2015, pp. 467–493. <https://doi.org/10.1017/jfm.2015.470>, URL <https://www.cambridge.org/core/product/0DAD57930F73C8158279B4064A49E9FF>.
- [25] Pope, S. B., *Turbulent Flows*, Cambridge University Press, 2000. <https://doi.org/10.1088/0957-0233/12/11/705>, URL <http://dx.doi.org/10.1088/0957-0233/12/11/705>.
- [26] Sørensen, J., Shen, W., and Munduate, X., “Analysis of wake states by a full-field actuator disc model,” *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, Vol. 1, No. 2, 1998, pp. 73–88.
- [27] Ferreira, C., Yu, W., Sala, A., and Viré, A., “Dynamic inflow model for a floating horizontal axis wind turbine in surge motion,” *Wind Energ. Sci.*, Vol. 7, No. 2, 2022, pp. 469–485. <https://doi.org/10.5194/wes-7-469-2022>, URL <https://wes.copernicus.org/articles/7/469/2022/>.
- [28] Li, Z., Liu, X., and Yang, X., “Review of Turbine Parameterization Models for Large-Eddy Simulation of Wind Turbine Wakes,” *Energies*, Vol. 15, No. 18, 2022, p. 6533. <https://doi.org/10.3390/en15186533>, URL <https://www.mdpi.com/1996-1073/15/18/6533>.
- [29] Glauert, H., *Airplane Propellers*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1935, pp. 169–360. https://doi.org/10.1007/978-3-642-91487-4_3.

- [30] Clark, D. R., and Leiper, A. C., "The free wake analysis," *Journal of the American Helicopter Society*, Vol. 15, No. 1, 1970, pp. 3–11.
- [31] McTavish, S., Feszty, D., and Nitzsche, F., "An experimental and computational assessment of blockage effects on wind turbine wake development," *Wind Energy*, Vol. 17, No. 10, 2014, pp. 1515–1529. <https://doi.org/10.1002/we.1648>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1648>.
- [32] Sørensen, J. N., and Myken, A., "Unsteady actuator disc model for horizontal axis wind turbines," *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 39, No. 1, 1992, pp. 139–149. [https://doi.org/10.1016/0167-6105\(92\)90540-Q](https://doi.org/10.1016/0167-6105(92)90540-Q), URL <https://www.sciencedirect.com/science/article/pii/016761059290540Q>.
- [33] Sørensen, J. N., and Shen, W. Z., "Numerical modeling of wind turbine wakes," *Journal of Fluids Engineering*, Vol. 124, No. 2, 2002, pp. 393–399.
- [34] Shen, W. Z., Sørensen, J. N., and Zhang, J., "Actuator surface model for wind turbine flow computations," *2007 European Wind Energy Conference and Exhibition*, 2007.
- [35] Shen, W. Z., Zhang, J. H., and Sørensen, J. N., "The Actuator Surface Model: A New Navier–Stokes Based Model for Rotor Computations (for VAWT)," *Journal of Solar Energy Engineering*, Vol. 131, No. 1, 2009. <https://doi.org/10.1115/1.3027502>.
- [36] Helmis, C. G., Whale, J., Papadopoulos, K. H., Anderson, C., Asimakopoulous, D., and Skyner, D., "A comparative laboratory and full-scale study of the near wake structure of a wind turbine," *Proceedings of the 5th European Wind Energy Association Conference*, 1994.
- [37] Simms, D., Hand, M., Jager, D., Cotrell, J., Robinson, M., Schreck, S., Larwood, S., and Fingersh, L., *Wind tunnel testing of NREL's unsteady aerodynamics experiment*, American Institute of Aeronautics & Astronautics, 2001. <https://doi.org/10.2514/6.2001-35>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2001-35>.
- [38] Dasari, T., "Near-wake studies of a utility-scale wind turbine using natural snowfall-based visualization and Particle Image Velocimetry," Phd thesis, University of Minnesota, 2018.
- [39] Wang, X., Cai, C., Cai, S.-G., Wang, T., Wang, Z., Song, J., Rong, X., and Li, Q., "A review of aerodynamic and wake characteristics of floating offshore wind turbines," *Renewable and Sustainable Energy Reviews*, Vol. 175, 2023, p. 113144. <https://doi.org/10.1016/j.rser.2022.113144>, URL <https://www.sciencedirect.com/science/article/pii/S1364032122010255>.
- [40] Rockel, S., Camp, E., Schmidt, J., Peinke, J., Cal, R. B., and Hölling, M., "Experimental Study on Influence of Pitch Motion on the Wake of a Floating Wind Turbine Model," *Energies*, Vol. 7, No. 4, 2014, pp. 1954–1985. URL <https://www.mdpi.com/1996-1073/7/4/1954>.
- [41] Rockel, S., Peinke, J., Hölling, M., and Cal, R. B., "Wake to wake interaction of floating wind turbine models in free pitch motion: An eddy viscosity and mixing length approach," *Renewable Energy*, Vol. 85, 2016, pp. 666–676. <https://doi.org/10.1016/j.renene.2015.07.012>, URL <https://www.sciencedirect.com/science/article/pii/S0960148115301087>.
- [42] Hu, H., Khosravi, M., and Sarkar, P., "An experimental investigation on the performance and the wake characteristics of a wind turbine subjected to surge motion," *33rd wind energy symposium*, 2015, p. 1207.
- [43] Bayati, I., Belloli, M., Facchinetti, A., and Giappino, S., "Wind Tunnel Tests on Floating Offshore Wind Turbines: A Proposal for Hardware-in-the-Loop Approach to Validate Numerical Codes," *Wind Engineering*, Vol. 37, No. 6, 2013, pp. 557–568. <https://doi.org/10.1260/0309-524x.37.6.557>, URL <https://journals.sagepub.com/doi/abs/10.1260/0309-524X.37.6.557>.

- [44] Bayati, I., Belloli, M., Bernini, L., and Zasso, A., “Wind Tunnel Wake Measurements of Floating Offshore Wind Turbines,” *Energy Procedia*, Vol. 137, 2017, pp. 214–222. <https://doi.org/10.1016/j.egypro.2017.10.375>, URL <https://www.sciencedirect.com/science/article/pii/S1876610217353584>.
- [45] Bayati, I., Facchinetti, A., Fontanella, A., Giberti, H., and Belloli, M., “A wind tunnel/HIL setup for integrated tests of Floating Offshore Wind Turbines,” *Journal of Physics: Conference Series*, Vol. 1037, No. 5, 2018, p. 052025. <https://doi.org/10.1088/1742-6596/1037/5/052025>, URL <https://dx.doi.org/10.1088/1742-6596/1037/5/052025>.
- [46] Bayati, I., Facchinetti, A., Fontanella, A., Taruffi, F., and Belloli, M., “Analysis of FOWT dynamics in 2-DOF hybrid HIL wind tunnel experiments,” *Ocean Engineering*, Vol. 195, 2020, p. 106717. <https://doi.org/10.1016/j.oceaneng.2019.106717>, URL <https://www.sciencedirect.com/science/article/pii/S0029801819308303>.
- [47] Belloli, M., Bayati, I., Facchinetti, A., Fontanella, A., Giberti, H., La Mura, F., Taruffi, F., and Zasso, A., “A hybrid methodology for wind tunnel testing of floating offshore wind turbines,” *Ocean Engineering*, Vol. 210, 2020, p. 107592. <https://doi.org/10.1016/j.oceaneng.2020.107592>, URL <https://www.sciencedirect.com/science/article/pii/S0029801820305953>.
- [48] Fontanella, A., Bayati, I., Mikkelsen, R., Belloli, M., and Zasso, A., “UNAFLOW: a holistic wind tunnel experiment about the aerodynamic response of floating wind turbines under imposed surge motion,” *Wind Energ. Sci.*, Vol. 6, No. 5, 2021, pp. 1169–1190. <https://doi.org/10.5194/wes-6-1169-2021>, URL <https://wes.copernicus.org/articles/6/1169/2021/>.
- [49] Fontanella, A., Zasso, A., and Belloli, M., “Wind tunnel investigation of the wake-flow response for a floating turbine subjected to surge motion,” *Journal of Physics: Conference Series*, Vol. 2265, No. 4, 2022, p. 042023. <https://doi.org/10.1088/1742-6596/2265/4/042023>, URL <https://dx.doi.org/10.1088/1742-6596/2265/4/042023>.
- [50] Hu, H., Khosravi, M., and Sarkar, P., “An experimental investigation on the aeromechanic performance and wake characteristics of a wind turbine model subjected to pitch motions,” *34th wind energy symposium*, 2016, p. 1997.
- [51] Fontanella, A., Da Pra, G., and Belloli, M., “Integrated Design and Experimental Validation of a Fixed-Pitch Rotor for Wind Tunnel Testing,” *Energies*, Vol. 16, No. 5, 2023, p. 2205. URL <https://www.mdpi.com/1996-1073/16/5/2205>.
- [52] Whale, J., Anderson, C. G., Bareiss, R., and Wagner, S., “An experimental and numerical study of the vortex structure in the wake of a wind turbine,” *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 84, No. 1, 2000, pp. 1–21. [https://doi.org/10.1016/S0167-6105\(98\)00201-3](https://doi.org/10.1016/S0167-6105(98)00201-3), URL <https://www.sciencedirect.com/science/article/pii/S0167610598002013>.
- [53] Chamorro, L. P., Arndt, R. E. A., and Sotiropoulos, F., “Reynolds number dependence of turbulence statistics in the wake of wind turbines,” *Wind Energy*, Vol. 15, No. 5, 2012, pp. 733–742. <https://doi.org/10.1002/we.501>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.501>.
- [54] McTavish, S., Feszty, D., and Nitzsche, F., “Evaluating Reynolds number effects in small-scale wind turbine experiments,” *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 120, 2013, pp. 81–90. <https://doi.org/10.1016/j.jweia.2013.07.006>, URL <https://www.sciencedirect.com/science/article/pii/S0167610513001372>.
- [55] de Vaal, J. B., Hansen, M. O. L., and Moan, T., “Effect of wind turbine surge motion on rotor thrust and induced velocity,” *Wind Energy*, Vol. 17, No. 1, 2014, pp. 105–121. <https://doi.org/10.1002/we.1562>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1562>.

- [56] Sant, T., Bonnici, D., Farrugia, R., and Micallef, D., “Measurements and modelling of the power performance of a model floating wind turbine under controlled conditions,” *Wind Energy*, Vol. 18, No. 5, 2015, pp. 811–834. <https://doi.org/10.1002/we.1730>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1730>.
- [57] Ramos-García, N., Kontos, S., Pegalajar-Jurado, A., González Horcas, S., and Bredmose, H., “Investigation of the floating IEA Wind 15 MW RWT using vortex methods Part I: Flow regimes and wake recovery,” *Wind Energy*, Vol. 25, No. 3, 2021, pp. 468–504. <https://doi.org/10.1002/we.2682>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2682>.
- [58] Fang, Y., Li, G., Duan, L., Han, Z., and Zhao, Y., “Effect of surge motion on rotor aerodynamics and wake characteristics of a floating horizontal-axis wind turbine,” *Energy*, Vol. 218, 2021, p. 119519. <https://doi.org/10.1016/j.energy.2020.119519>, URL <https://www.sciencedirect.com/science/article/pii/S0360544220326268>.
- [59] Li, Y., Yu, W., and Sarlak, H., “Wake Structures and Performance of Wind Turbine Rotor With Harmonic Surging Motions Under Laminar and Turbulent Inflows,” *Wind Energy*, Vol. 27, No. 12, 2024, pp. 1499–1525. <https://doi.org/10.1002/we.2949>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2949>.
- [60] Micallef, D., and Sant, T., “Loading effects on floating offshore horizontal axis wind turbines in surge motion,” *Renewable Energy*, Vol. 83, 2015, pp. 737–748. <https://doi.org/10.1016/j.renene.2015.05.016>, URL <https://www.sciencedirect.com/science/article/pii/S0960148115003936>.
- [61] Farrugia, R., Sant, T., and Micallef, D., “A study on the aerodynamics of a floating wind turbine rotor,” *Renewable Energy*, Vol. 86, 2016, pp. 770–784. <https://doi.org/10.1016/j.renene.2015.08.063>, URL <https://www.sciencedirect.com/science/article/pii/S0960148115302639>.
- [62] Cormier, M., Caboni, M., Lutz, T., Boorsma, K., and Krämer, E., “Numerical analysis of unsteady aerodynamics of floating offshore wind turbines,” *Journal of Physics: Conference Series*, Vol. 1037, No. 7, 2018, p. 072048. <https://doi.org/10.1088/1742-6596/1037/7/072048>, URL <https://dx.doi.org/10.1088/1742-6596/1037/7/072048>.
- [63] Shen, X., Chen, J., Hu, P., Zhu, X., and Du, Z., “Study of the unsteady aerodynamics of floating wind turbines,” *Energy*, Vol. 145, 2018, pp. 793–809. <https://doi.org/10.1016/j.energy.2017.12.100>, URL <https://www.sciencedirect.com/science/article/pii/S0360544217321436>.
- [64] Lienard, C., Boisard, R., and Daudin, C., “Aerodynamic behavior of a floating offshore wind turbine,” *AIAA Journal*, Vol. 58, No. 9, 2020, pp. 3835–3847.
- [65] Kyle, R., Lee, Y. C., and Früh, W.-G., “Propeller and vortex ring state for floating offshore wind turbines during surge,” *Renewable Energy*, Vol. 155, 2020, pp. 645–657. <https://doi.org/10.1016/j.renene.2020.03.105>, URL <https://www.sciencedirect.com/science/article/pii/S0960148120304377>.
- [66] Dong, J., and Viré, A., “Comparative analysis of different criteria for the prediction of vortex ring state of floating offshore wind turbines,” *Renewable Energy*, Vol. 163, 2021, pp. 882–909. <https://doi.org/10.1016/j.renene.2020.08.027>, URL <https://www.sciencedirect.com/science/article/pii/S0960148120312647>.
- [67] Arabgolarcheh, A., Jannesarahmadi, S., and Benini, E., “Modeling of near wake characteristics in floating offshore wind turbines using an actuator line method,” *Renewable Energy*, Vol. 185, 2022, pp. 871–887. <https://doi.org/10.1016/j.renene.2021.12.099>, URL <https://www.sciencedirect.com/science/article/pii/S0960148121018267>.
- [68] Kyle, R., and Früh, W.-G., “The transitional states of a floating wind turbine during high levels of surge,” *Renewable Energy*, Vol. 200, 2022, pp. 1469–1489. <https://doi.org/10.1016/j.renene.2022.10.034>, URL <https://www.sciencedirect.com/science/article/pii/S0960148122015270>.

- [69] Fu, S., Jin, Y., Zheng, Y., and Chamorro, L. P., “Wake and power fluctuations of a model wind turbine subjected to pitch and roll oscillations,” *Applied Energy*, Vol. 253, 2019, p. 113605. <https://doi.org/10.1016/j.apenergy.2019.113605>, URL <https://www.sciencedirect.com/science/article/pii/S0306261919312796>.
- [70] Kleine, V. G., Franceschini, L., Carmo, B. S., Hanifi, A., and Henningson, D. S., “The stability of wakes of floating wind turbines,” *Physics of Fluids*, Vol. 34, No. 7, 2022. <https://doi.org/10.1063/5.0092267>, URL 10.1063/5.0092267.
- [71] Messmer, T., Hölling, M., and Peinke, J., “Enhanced recovery caused by nonlinear dynamics in the wake of a floating offshore wind turbine,” *Journal of Fluid Mechanics*, Vol. 984, 2024, p. A66. <https://doi.org/10.1017/jfm.2024.175>, URL <https://www.cambridge.org/core/product/70424633A9AECB7771F0AEF252F341EF>.
- [72] Wiegant, E., and Verzijlbergh, R., “GRASP model description & validation report,” , 2019.
- [73] De Tavernier, D., and von Terzi, D., “The emergence of supersonic flow on wind turbines,” *Journal of Physics: Conference Series*, Vol. 2265, No. 4, 2022, p. 042068. <https://doi.org/10.1088/1742-6596/2265/4/042068>, URL <https://dx.doi.org/10.1088/1742-6596/2265/4/042068>.
- [74] Smagorinsky, J., “GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS: I. THE BASIC EXPERIMENT,” *Monthly Weather Review*, Vol. 91, No. 3, 1963, pp. 99–164. [https://doi.org/10.1175/1520-0493\(1963\)091<0099:GCEWTP>2.3.CO;2](https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2), URL https://journals.ametsoc.org/view/journals/mwre/91/3/1520-0493_1963_091_0099_gcewtp_2_3_co_2.xml.
- [75] Sullivan, P. P., McWilliams, J. C., and Moeng, C.-H., “A subgrid-scale model for large-eddy simulation of planetary boundary-layer flows,” *Boundary-Layer Meteorology*, Vol. 71, No. 3, 1994, pp. 247–276. <https://doi.org/10.1007/BF00713741>.
- [76] Vreman, A. W., “An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications,” *Physics of Fluids*, Vol. 16, No. 10, 2004, pp. 3670–3681. <https://doi.org/10.1063/1.1785131>, URL 10.1063/1.1785131.
- [77] Verstappen, R., Bose, S., Lee, J., Choi, H., and Moin, P., “A dynamic eddy-viscosity model based on the invariants of the rate-of-strain,” *Proceedings of the Summer Program*, 2010, pp. 183–192.
- [78] Rozema, W., Bae, H. J., Moin, P., and Verstappen, R., “Minimum-dissipation models for large-eddy simulation,” *Physics of Fluids*, Vol. 27, No. 8, 2015. <https://doi.org/10.1063/1.4928700>, URL 10.1063/1.4928700.
- [79] Martínez-Tossas, L. A., Churchfield, M. J., and Leonardi, S., “Large eddy simulations of the flow past wind turbines: actuator line and disk modeling,” *Wind Energy*, Vol. 18, No. 6, 2015. <https://doi.org/10.1002/we.1747>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1747>.
- [80] Verstappen, R., “When Does Eddy Viscosity Damp Subfilter Scales Sufficiently?” *Journal of Scientific Computing*, Vol. 49, No. 1, 2011, pp. 94–110. <https://doi.org/10.1007/s10915-011-9504-4>.
- [81] Cengel, Y., and Cimbala, J., *Fluid mechanics fundamentals and applications (si units)*, McGraw Hill, 2013.
- [82] Boyce, W. E., DiPrima, R. C., and Haines, C. W., *Elementary differential equations and boundary value problems*, Wiley New York, 1969.
- [83] Jha, P. K., Churchfield, M. J., Moriarty, P. J., and Schmitz, S., “Guidelines for Volume Force Distributions Within Actuator Line Modeling of Wind Turbines on Large-Eddy Simulation-Type Grids,” *Journal of Solar Energy Engineering*, Vol. 136, No. 3, 2014. <https://doi.org/10.1115/1.4026252>.

- [84] Churchfield, M. J., Schreck, S. J., Martinez, L. A., Meneveau, C., and Spalart, P. R., *An Advanced Actuator Line Method for Wind Energy Applications and Beyond*, AIAA, 2017. <https://doi.org/10.2514/6.2017-1998>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2017-1998>.
- [85] Martínez-Tossas, L. A., Churchfield, M. J., and Meneveau, C., “Optimal smoothing length scale for actuator line models of wind turbine blades based on Gaussian body force distribution,” *Wind Energy*, Vol. 20, No. 6, 2017.
- [86] Martínez-Tossas, L. A., and Meneveau, C., “Filtered lifting line theory and application to the actuator line model,” *Journal of Fluid Mechanics*, Vol. 863, 2019. <https://doi.org/10.1017/jfm.2018.994>, URL <https://www.cambridge.org/core/product/EA90C34B7D48030D725B3AFDC06501DF>.
- [87] Stanly, R., Martínez-Tossas, L. A., Frankel, S. H., and Delorme, Y., “Large-eddy simulation of a wind turbine using a filtered actuator line model,” *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 222, 2022, p. 104868.
- [88] Churchfield, M. J., Lee, S., Schmitz, S., and Wang, Z., “Modeling wind turbine tower and nacelle effects within an actuator line model,” *33rd Wind Energy Symposium*, 2015, p. 214.
- [89] Delft High Performance Computing Centre (DHPC), “DelftBlue Supercomputer (Phase 2),” , 2024. URL <https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2>.

A

Simulation tuning

A.1. Experimental results




Figure A.1: Normalised time-averaged streamwise velocity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.2: Normalised time-averaged out-of-plane vorticity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.3: Time-averaged TI field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.4: Normalised time-averaged streamwise-vertical Reynolds stress field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.5: Normalised time-averaged vertical velocity field from the experiments for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.

A.2. Base simulation




Figure A.6: Normalised time-averaged streamwise velocity field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.7: Normalised time-averaged streamwise velocity field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.8: Time-averaged TI field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.9: Normalised time-averaged streamwise-vertical Reynolds stress field from the base simulation for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.

A.3. Turbulent inflow




Figure A.10: Normalised time-averaged streamwise velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.11: Normalised time-averaged streamwise velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.12: Time-averaged TI field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.13: Normalised time-averaged streamwise-vertical Reynolds stress field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.




Figure A.14: Normalised time-averaged vertical velocity field from the simulations with turbulent inflow for, from top to bottom, the stationary, Low-Frequency Surge, High-Frequency Surge, Low-Frequency Pitch and High-Frequency Pitch cases.

A.4. FALM




Figure A.15: Normalised time-averaged streamwise velocity difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.




Figure A.16: Time-averaged TI difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.




Figure A.17: Normalised time-averaged streamwise-vertical Reynolds stress difference between simulations using the FALM and experiments (solid line). Additionally, the difference between the simulations using the ALM and experiments are presented (dotted line). All profiles are along the vertical direction of the domain. The presented profiles are at x/D equal to 0.7, 1.9, 3.1 and 4.4.

A.5. Including Nacelle

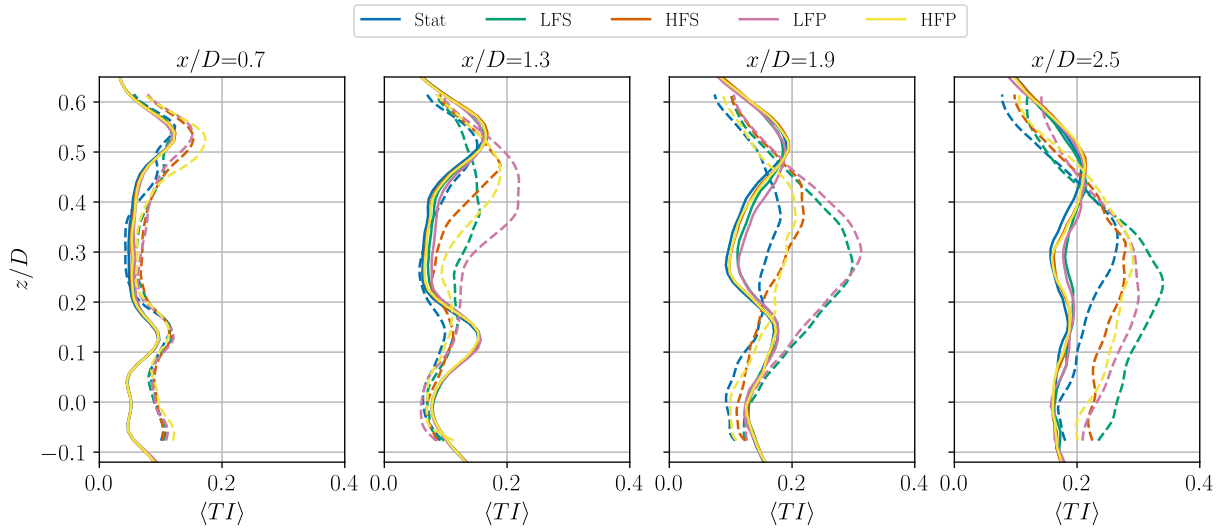


Figure A.18: Time-averaged TI along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5.

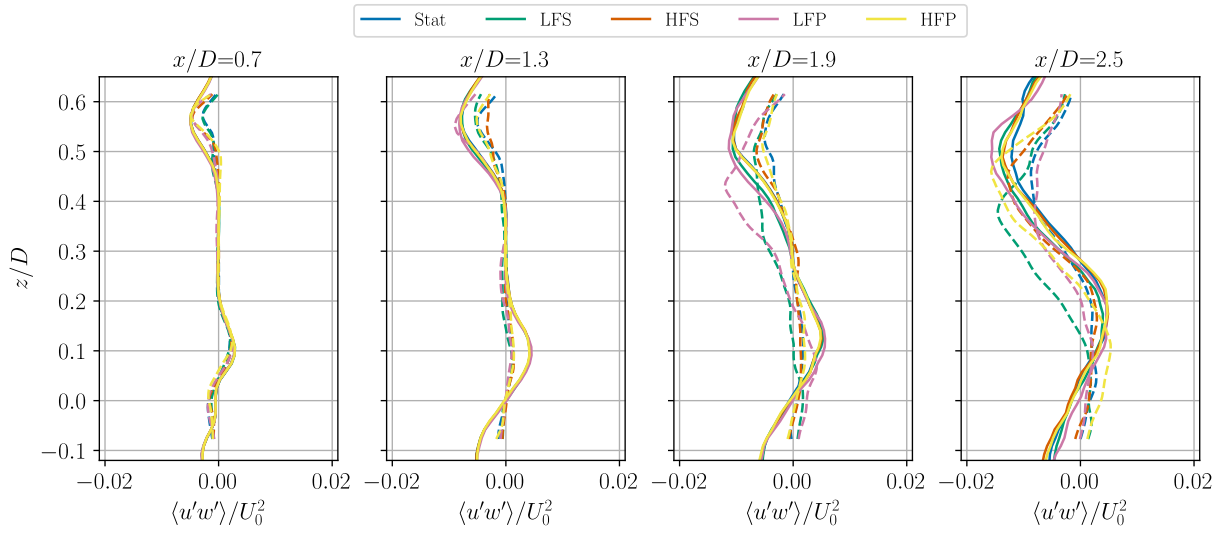


Figure A.19: Normalised time-averaged out-of-plane velocity along the vertical direction from the simulations including a nacelle for all motion cases. The presented profiles are at x/D equal to . 0.7, 1.3, 1.9 and 2.5.

A.6. High resolution

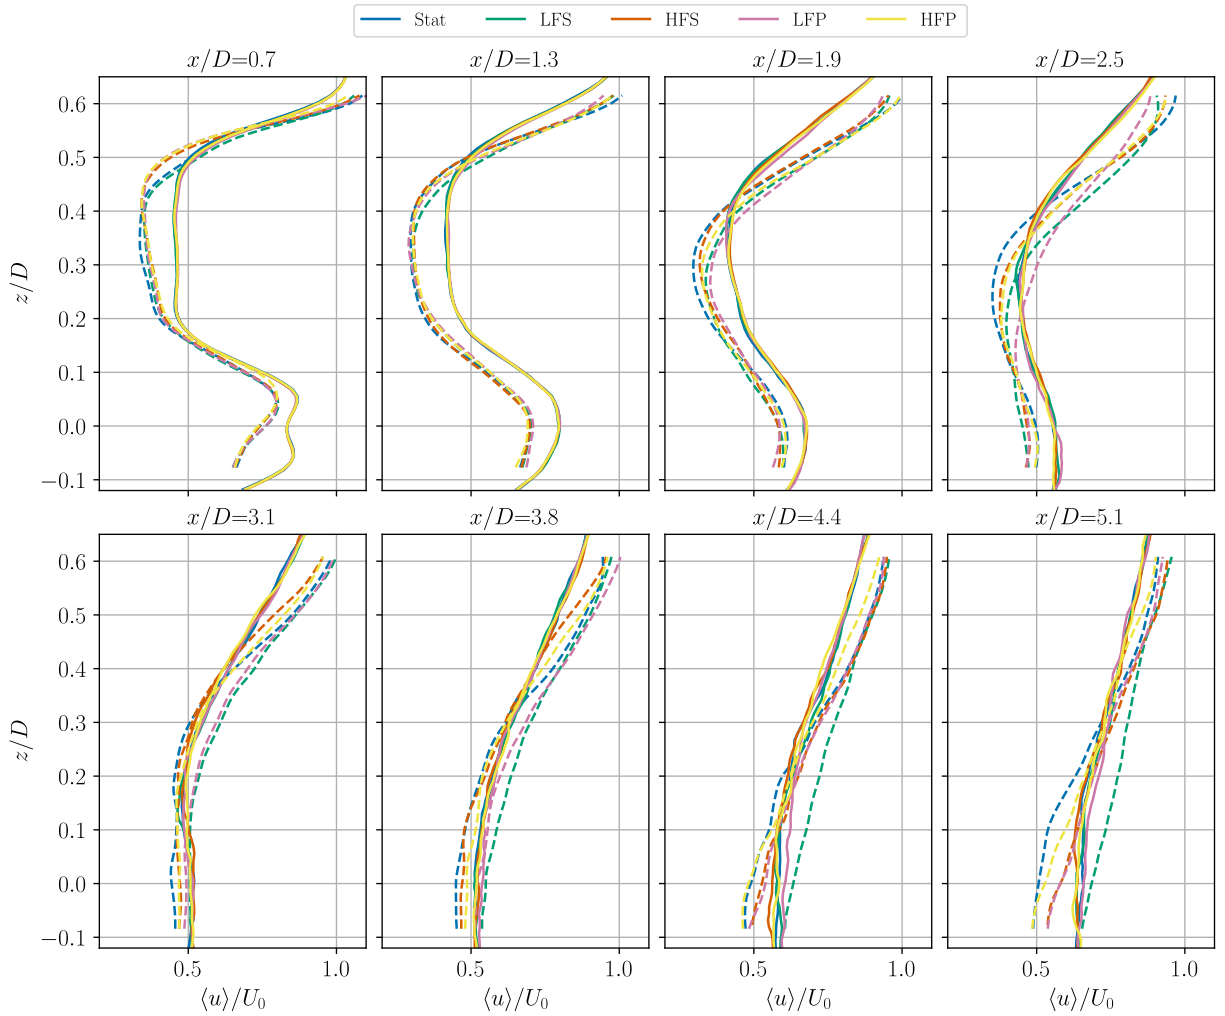


Figure A.20: Normalised time-averaged streamwise velocity along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

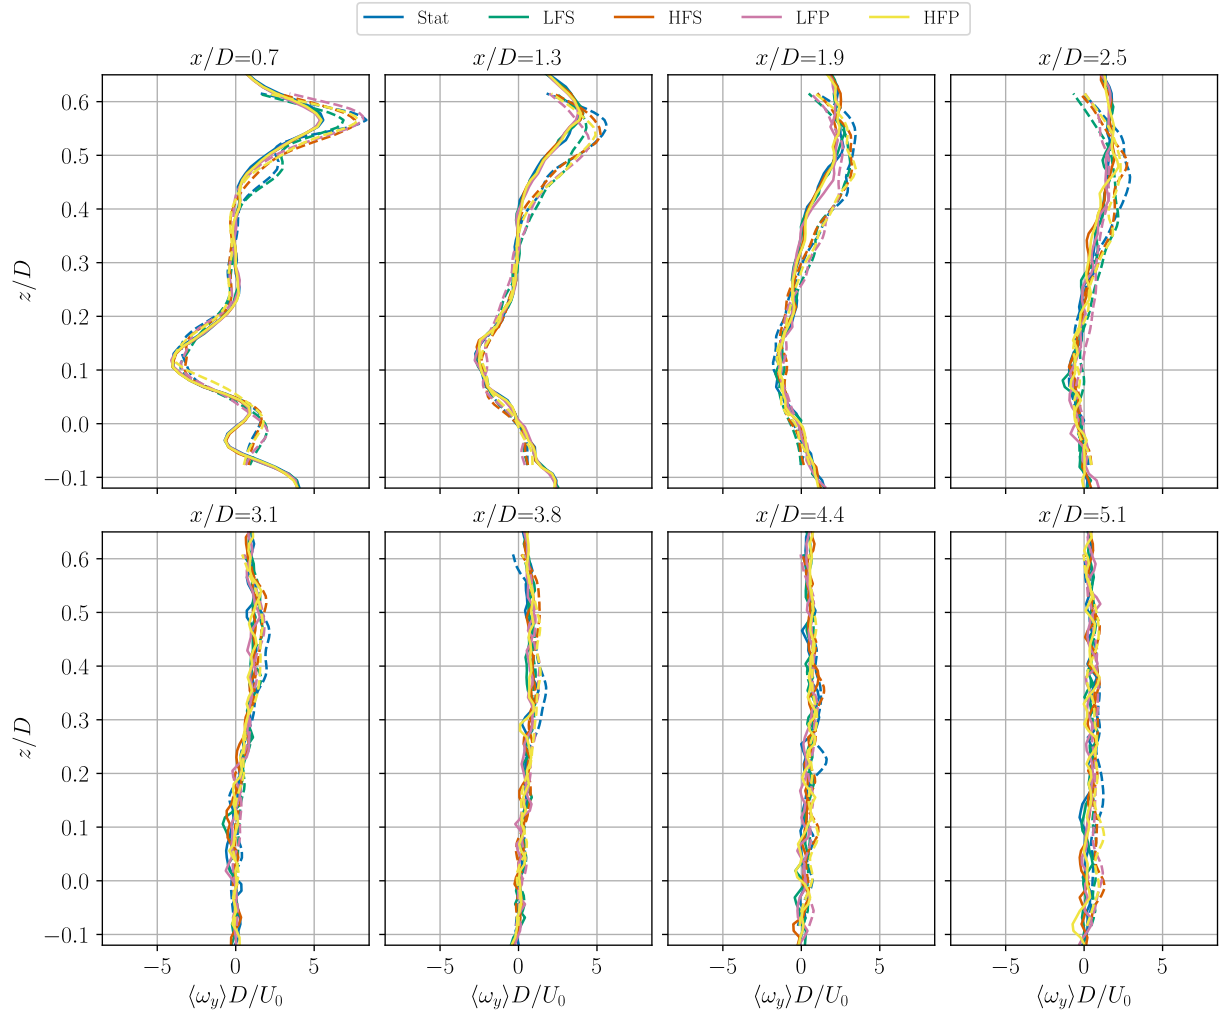


Figure A.21: Normalised time-averaged out-of-plane vorticity along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

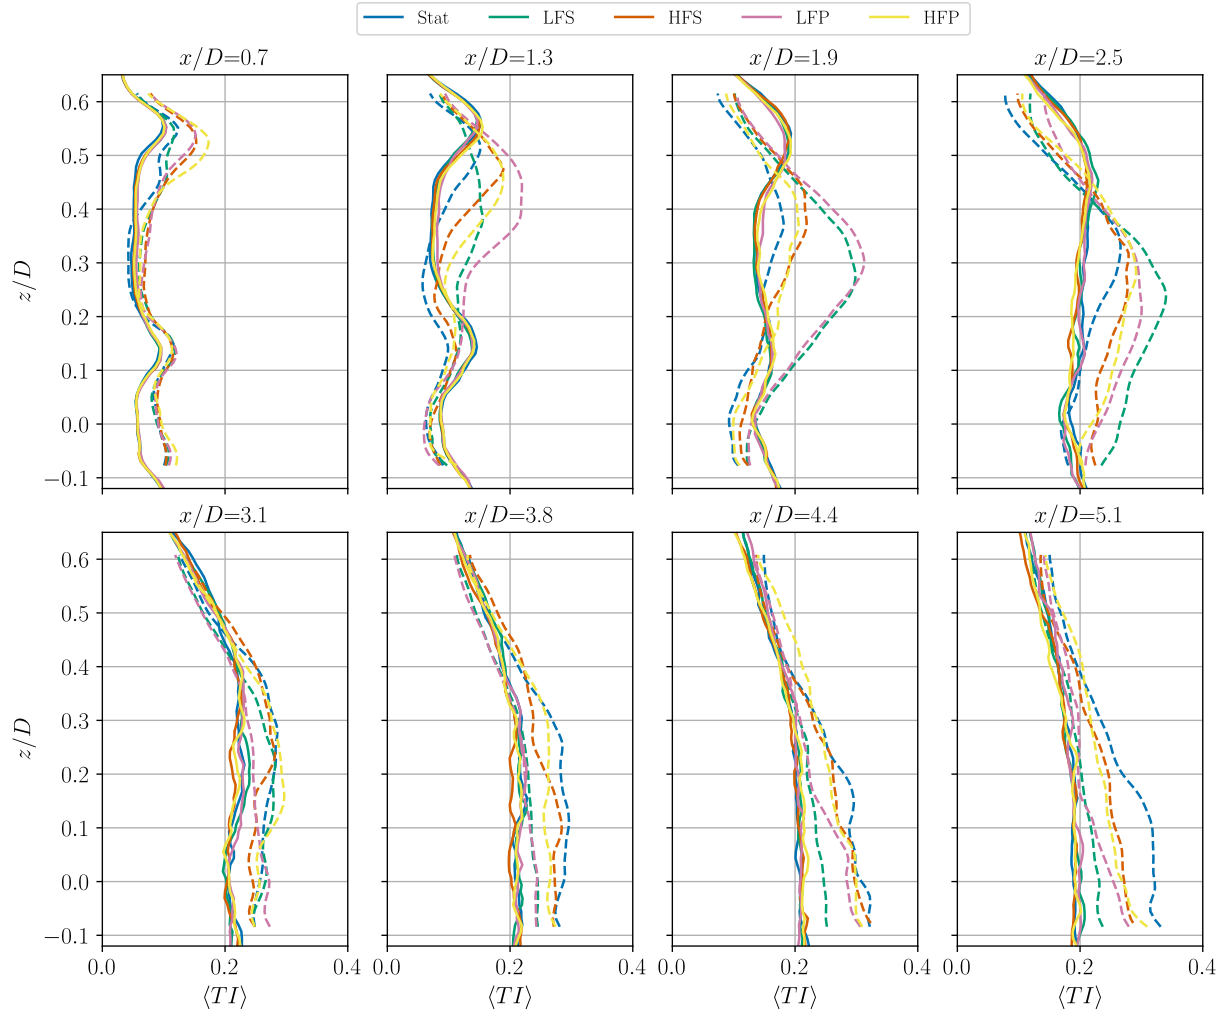


Figure A.22: Time-averaged TI along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

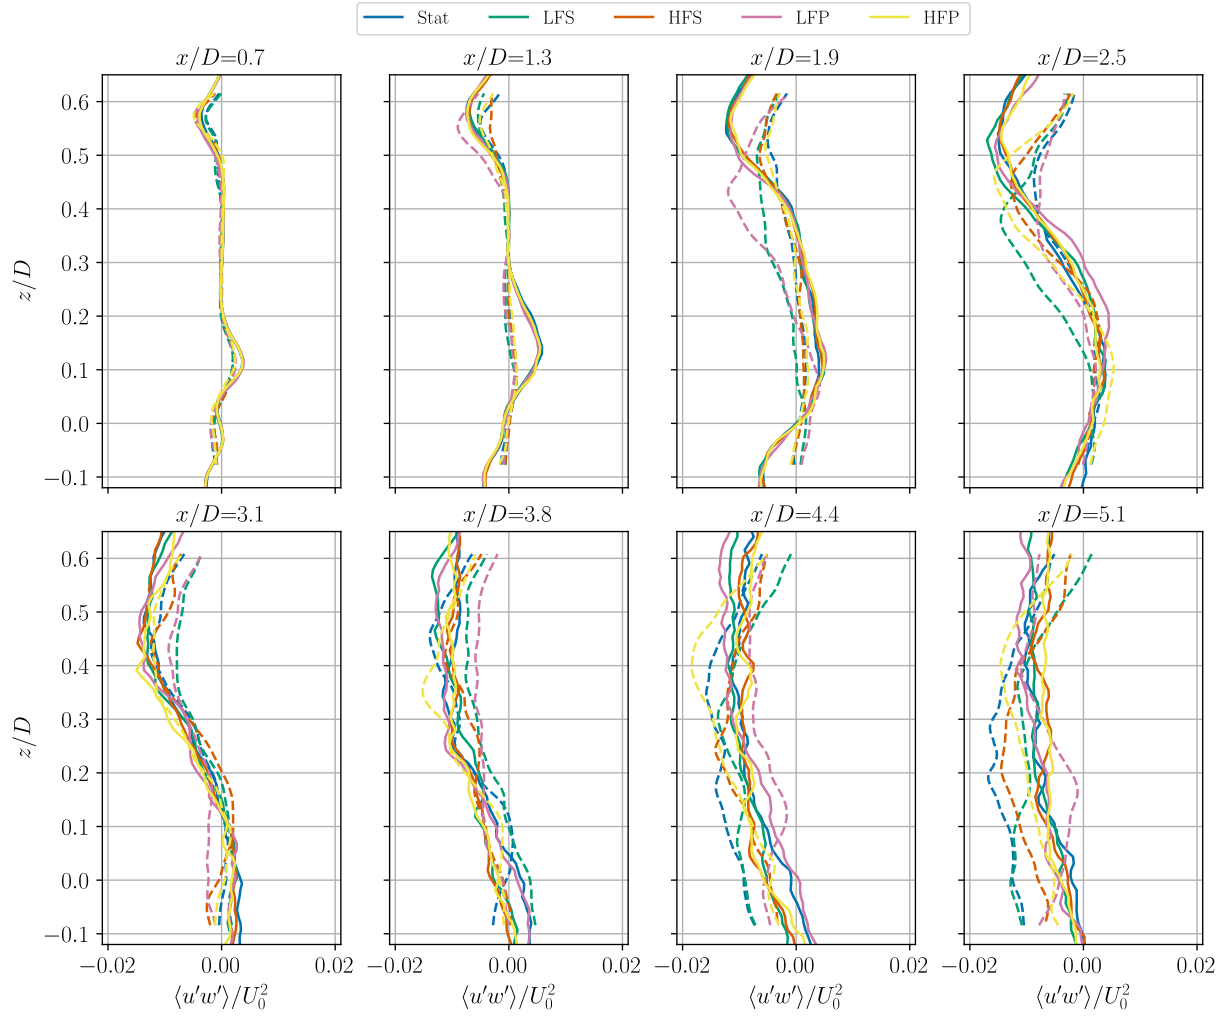


Figure A.23: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the high-resolution simulations with turbulent inflow (solid lines) and experiments (dashed lines). The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1.

B

Sub-Grid Scale Model and Advection Scheme Choice

In this work the Smagorinsky model has been used to model the unresolved turbulence scales since it has been used widely in actuator line modelling research [14, 7, 59]. However, GRASP also contains other models as discussed in Chapter 4. Especially, the Verstappen model is interesting as an improvement on the use of the Smagorinsky model. Besides to this model, the Rozema model would be of interest if a non-isotropic mesh was used.

During the early stage of this study, the parameter analysis assessed the effect of the Verstappen model compared to the Smagorinsky model. Likewise, the effect of the different orders of advection schemes was analysed. However, negligible differences were observed between the models and advection schemes in the normalised time-averaged streamwise velocity profiles at 2 rotor diameters behind the rotor plane. As will be presented in this appendix, this observation was correct. However, adding turbulence to the inflow and letting the near wake transition towards the far wake changes the conclusions. Unfortunately, this observation was made at a later stages in this work such that adopting a new SGS model and again running the investigated cases was not feasible.

This appendix will assess the differences between the Smagorinsky and Verstappen models; the Verstappen model will be combined with the second and fifth order advection schemes with the appropriate constants as discussed in Section 4.1.6. The use of two advection schemes presents to opportunity to compare them assess the effects. In literature, the fourth order scheme is mentioned for the Verstappen model with an appropriate constant. However, this order is not implemented in GRASP so a fifth order scheme has been used instead with the same constant of 0.486. The second order advection scheme Verstappen model uses a constant of 0.333.

The results will be presented as follows. First, the turbulence spectrum in the tip vortex region will be presented in Section B.1 to show a non-negligible difference in turbulence modelling exists. Secondly, the normalised time-averaged velocity will be presented in Section B.2. Lastly, the streamwise-vertical Reynolds stresses are presented in Section B.3 since they show the effect of the different SGS models and advection schemes best. It is also representative for the effects which are observed in the out-of-plane vorticity and TI field.

B.1. Turbulence Spectrum

The most turbulence is present in the tip region where the shear stresses are high. Therefore, the spatial turbulence spectrum will be taken from a plane with its lower corner at $(x/D = 0.6, y/D = 0, z/D = 0.4)$ and the upper corner at $(x/D = 0.6, y/D = 0, z/D = 0.6)$. The spectra for the different models are presented in Figure B.1.

In these spectra, the energy-containing range can be seen at wavelength around 0.1 m^{-1} . This is followed by a long inertial subrange, and ends with at the dissipation range. At a wave number of 2 m^{-1} there starts to be a significant difference between the Smagorinsky model and the Verstappen model. It appears that there is more energy at the start of the dissipation range when the Smagorinsky model is used. This could indicate a cusp, which is described in [25] for under-resolved grids. It would mean that the Smagorinsky model does not model enough dissipation. Therefore, the Verstappen model is more appropriate. Additionally, no difference between the two advection schemes is observed.




Figure B.1: Spatial turbulence spectra for the turbulent inflow simulation setup (see Chapter 8) with the Smagorinsky SGS model and the Verstappen model with a second and fifth order advection scheme. The grey line indicates the $-5/3$ slope in the inertial subrange.

B.2. Streamwise Velocity

This section discusses the normalised streamwise velocity profiles for the different SGS models and advection schemes at eight distances behind the rotor plane inside the wake.

Figure B.2 presents these profiles, together with the profiles from the experiment of the stationary turbine. Remember that the tip vortex breakdown occurs around $x/D = 1.6$. After this point the vorticity diffuses fast due to added random turbulence. It can be seen that there are minimal differences inside the wake's streamwise velocity profiles until $x/D = 1.9$. From $x/D = 2.5$ onwards, there are some small differences between the wake profiles. It appears that the profiles fluctuate around each other with the Verstappen model with the fifth order advection scheme to be the smoothest. There is no part of the far wake which is modelled better due to the use of a particular model.

B.3. Reynolds stress

In Figure B.3, the normalised time-averaged Reynolds stresses have been presented at eight locations downwind of the rotor plane. In addition, the profiles of the experiment with the stationary wind turbine has been added as a reference. It can be seen that the differences in these profiles are more pronounced. In general, until $x/D = 1.3$ is affected little by the different SGS models and advection schemes. However, in the stress maximum due to the root vortex at $z/D = 0.1$ the Verstappen model causes a small decrease in stress compared to the Smagorinsky. Then, using the fifth order advection scheme causes a further decrease. This shows to be in the correct direction since the experiments show a far lower minimum. At $x/D = 1.9$, it appears again like the Smagorinsky and second order Verstappen model fluctuate around the fifth order Verstappen. However, there is no clear improvement compared to the experiment. This remains the case until the very end of the domain. However, it becomes clearer that the Verstappen model causes a decrease in fluctuations of the stresses, and that the increased order of the advection scheme further smooths the profiles. Unfortunately, this does not produce results that are closer to the experiment results. However, using the fifth order Verstappen model would improve the ability to analyse the far wake by reducing the high frequency fluctuations in the profiles.

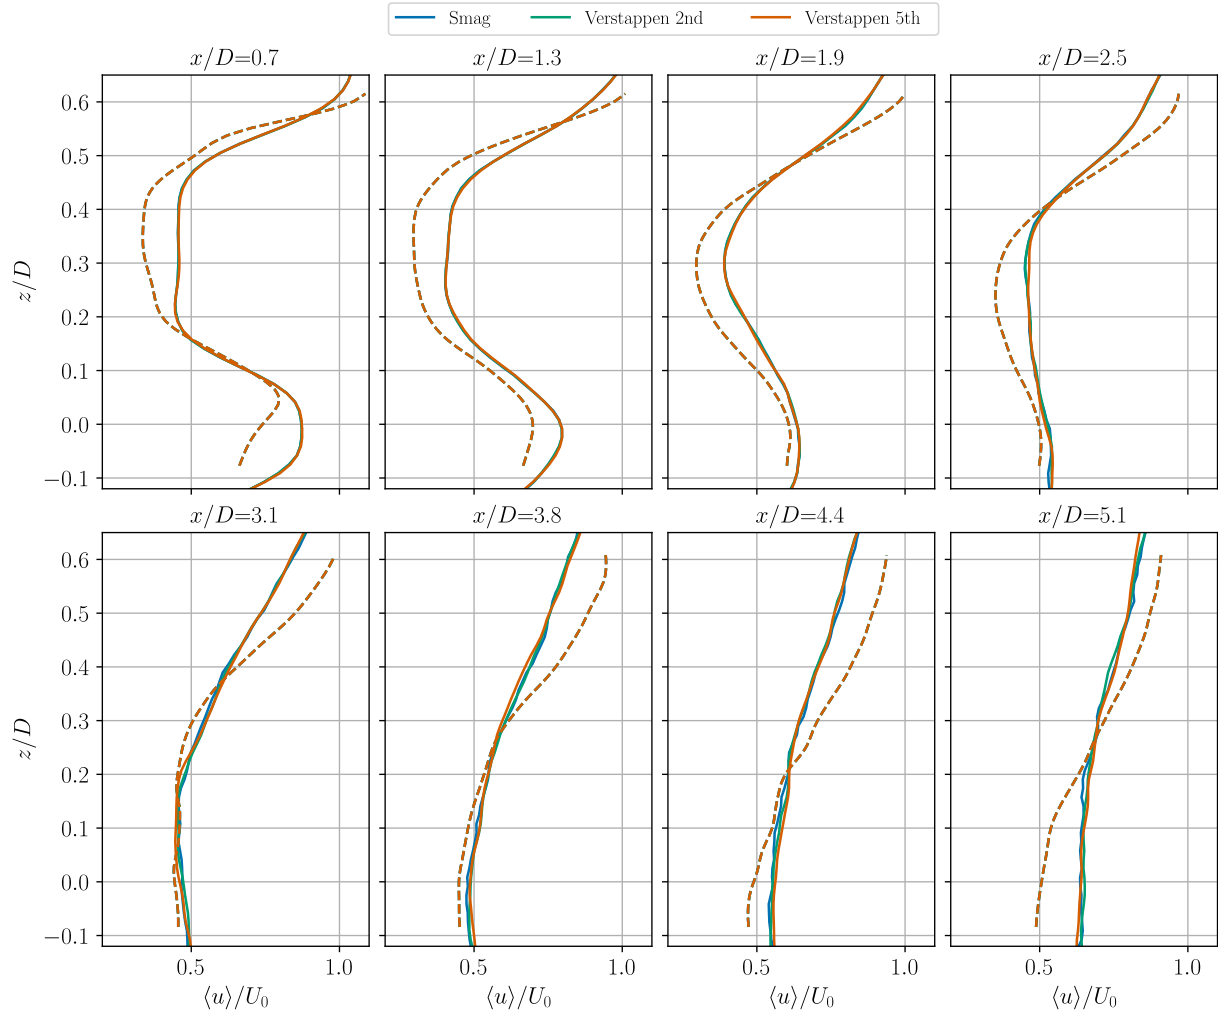


Figure B.2: Normalised time-averaged streamwise velocity along the vertical direction from the simulation with turbulent inflow. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. In addition, the experimental results for the stationary case have been added as the grey dashed line.

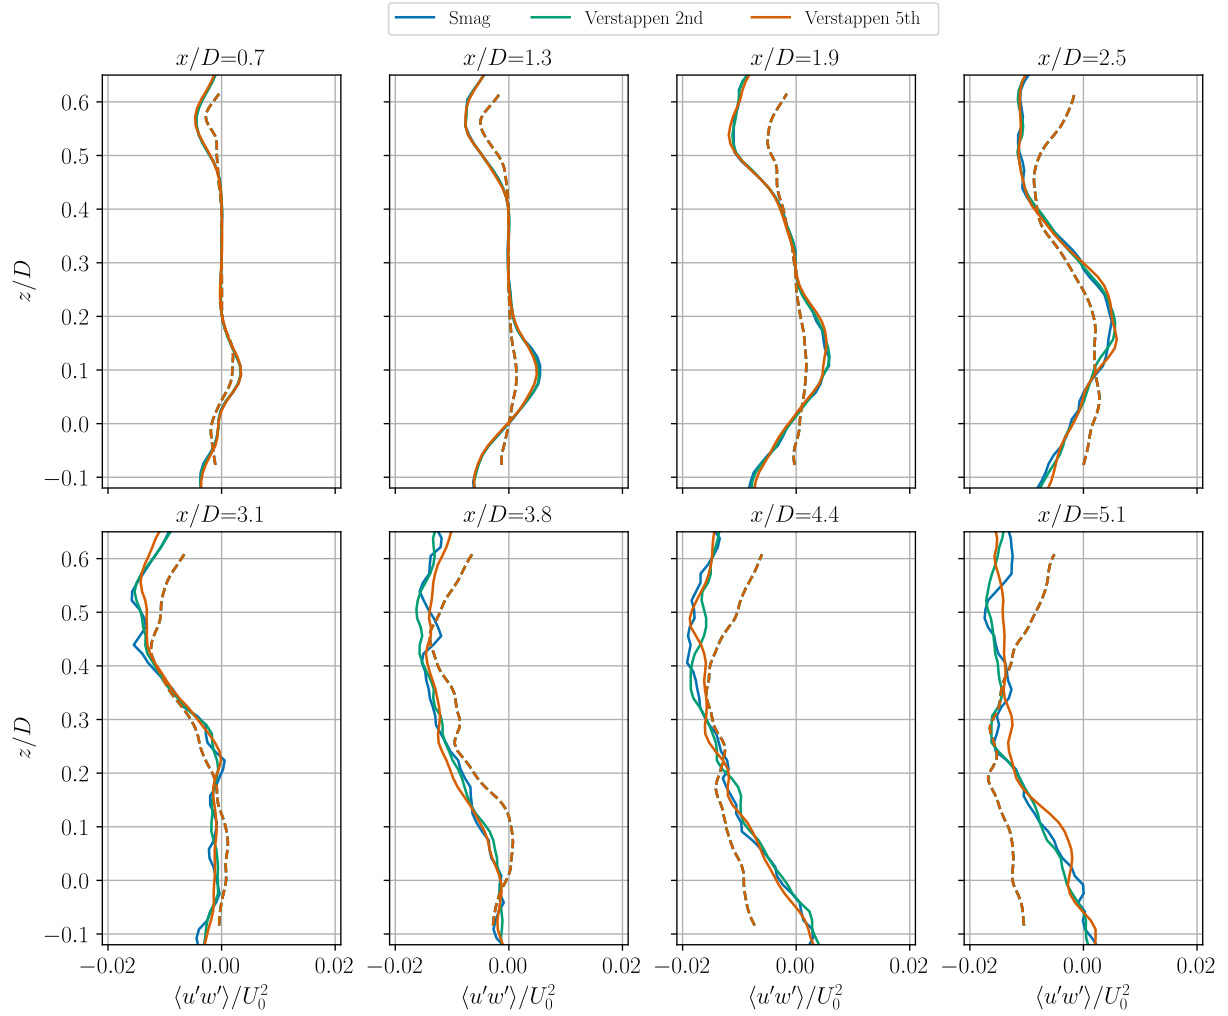


Figure B.3: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the simulation with turbulent inflow. The presented profiles are at x/D equal to 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4 and 5.1. In addition, the experimental results for the stationary case have been added as the grey dashed line.

Tower Modelling

In Section 9.2, the results of including the nacelle was discussed. It was shown that it can increase the accuracy of the near wake until $x/D = 1.3$. The next logical step is using the option in AspFAST to include the tower to model the drag of it on the flow using the single actuator drag line as discussed in Section 4.3.3. The tower diameter is equal to 8 cm; the drag coefficient is equal to 0.5 (for a cylinder) and 60 actuator points are used (no difference with 40 actuator points was observed). Additionally, the velocity has been measured half way between the rotor plane and the tower shell. Unfortunately, the results of the performed simulations which include the tower presented odd results and there was no time to find the problem. Therefore, the results have been left out of the main content of this work and are presented in this appendix such that it could potentially be solved in the future. Only the stationary results will be presented for clarity.

The effect of the current implementation of the tower has been presented for streamwise velocity, out-of-plane vorticity, TI and streamwise-vertical Reynolds stress in Section C.1, Section C.2, Section C.3 and Section C.4, respectively. Each section contains three simulations and the experimental results from the experiments (grey-dashed line). One simulation has already been presented in Section 9.1 where only the rotor is modelled and the FALM is used. Secondly, the low-resolution simulation including the nacelle has been presented in Section 9.2. The third simulation is including the tower. The profiles at x/D equal to 0.7, 1.3, 1.9 and 2.5 have been presented since this is enough to explain the problem.

C.1. Streamwise Velocity

In Figure C.1, the profiles for the normalised time-averaged velocity are presented. It can be seen that in the root region at $x/D = 0.7$, the velocity profile is more similar to the experiment than the other simulations. The drag of the tower is however not enough to reduce the velocity of the flow in the wake of the tower to an equal velocity as seen in the experiment. This is due to not modelling the tower as a body but as a line which does not stop flow from entering the volume. Besides this improvement, the other parts of the wake are significantly deteriorated by adding the tower. The root vortex velocity gradient has become less steep. Also the tip vortex velocity gradient has become less. Due to this, the wake core velocity has a significantly increase velocity. At $x/D = 1.3$, the improvement of in the root region is gone. However, the gradient of the velocity below the wake axis looks better compared to the simulation with only the rotor and nacelle. Therefore, it is believed that the overall higher velocity that is observed in the profiles is the cause of the lower accuracy in the root region. In the next profiles, the velocity profiles become more and more curves with a constant offset from each other.

C.2. Vorticity

In Figure C.2, the vorticity profiles are presented. Similar observation can be made as in the streamwise velocity results. At $x/D = 0.7$, the vorticity is improved in the root region due to the inclusion of the tower. The lower root vortex decreases significantly in size by including the tower. This is also the case at $x/D = 1.3$. However, in the root and tip vortices there is significantly less vorticity throughout the complete near wake. This explains where the velocity increase comes from as less velocity is induced by the vorticity. The wake core remains to contain zero vorticity and its height does not evolve differently from the other simulations throughout the profiles that are presented.

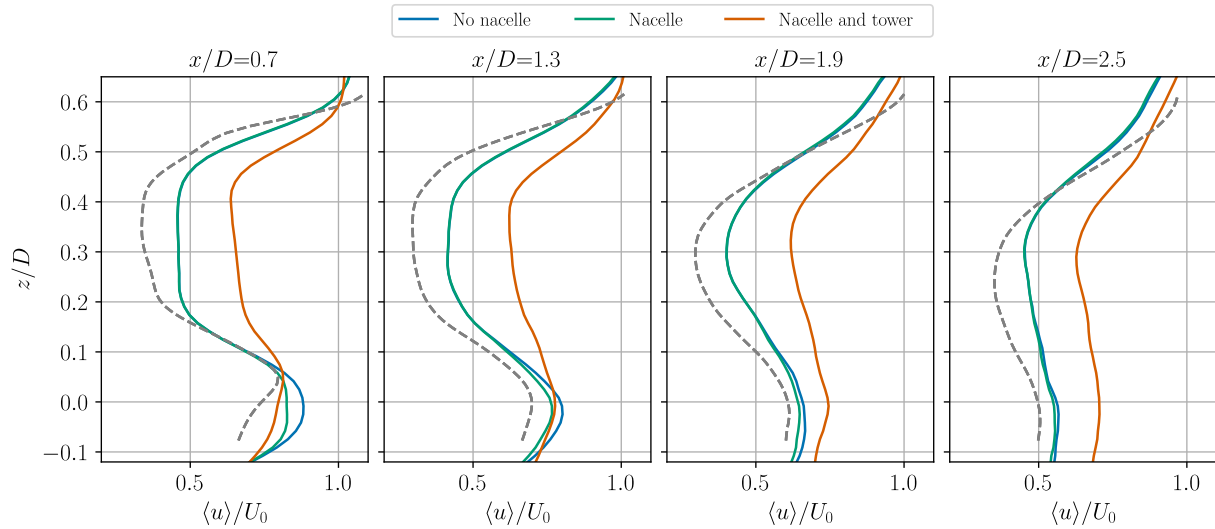


Figure C.1: Normalised time-averaged streamwise velocity along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.

C.3. Turbulence Intensity

The TI profiles are presented in Figure C.3. It shows lower TI values everywhere, except around the wake axis at $x/D = 0.7$. This shows that the nacelle still results in the same amount of turbulence. Besides, the tower results in a TI decrease around $z/D = -0.1$. Such a decrease can also be seen in the experiments, let it be at a higher TI. In the root and tip vortices there is especially a decrease in turbulence compared to the other simulations. Moreover, the wake core shows a small TI decrease due to adding the tower. Further downwind, the TI remains underpredicted and similar to the velocity results, the profiles become more and more at an constant offset from each other.

C.4. Reynolds stress

In Figure C.4, the normalised time-averaged streamwise-vertical Reynolds stresses are presented. Interestingly, the tower seems to improve the modelling of the stresses by lowering them. Below $z/D = -0.05$, the stress was too negative at $x/D = 0.7$. Including the tower makes the stress less negative. Therefore, less turbulence is generated in this region as was seen before. Downwind this persists such that at $x/D = 2.5$, there is a less negative stress below $z/D = 0.08$. Also, the stress becomes less positive in the root vortex region due to the lower velocity gradient. In the same way, the stress becomes less negative in the tip vortex due to the lower velocity gradient. The stress in the root vortex remains close to that observed in the experiments. However, After the tip vortex breakdown, the stress minimum starts to move downwards. As seen before this is not seen in the simulation. However, the stress minimum at $x/D = 2.5$ is almost equal in value as observed in the experiment.




Figure C.2: Normalised time-averaged out-of-plane vorticity along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.




Figure C.3: Time-averaged TI along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.




Figure C.4: Normalised time-averaged streamwise-vertical Reynolds stress along the vertical direction from the high-resolution simulations with turbulent inflow for all motion cases. The presented profiles are at x/D equal to 0.7, 1.3, 1.9 and 2.5. As a reference, the experimental result of the stationary motion case is included as the grey dashed line.

D

Vertical Shift Correction

As discussed in Chapter 8 and Chapter 10, a mismatch in the vertical position of the wake between the simulations and the experiments is observed, caused by differences in vertical velocity. This chapter outlines a simple method to correct for this mismatch. The approach for determining the required vertical shift based on vertical velocity differences is presented in Section D.1, followed by its application to the streamwise velocity results from Chapter 10 in Section D.2.

D.1. Correction Method

The procedure is illustrated in the block diagram in Figure D.1. At each streamwise location, the vertical velocities from both the experiments and the simulations are averaged along the vertical direction. The difference between these averaged values of the vertical velocities is then calculated.

The local convection time for the flow to travel a streamwise distance δx is determined from the vertically averaged local streamwise velocity. Multiplying this time by the vertical velocity difference gives the additional vertical position correction required over δx . By taking the cumulative sum, accounting for the previously applied vertical shifts, the total vertical correction at each streamwise location can then be determined.

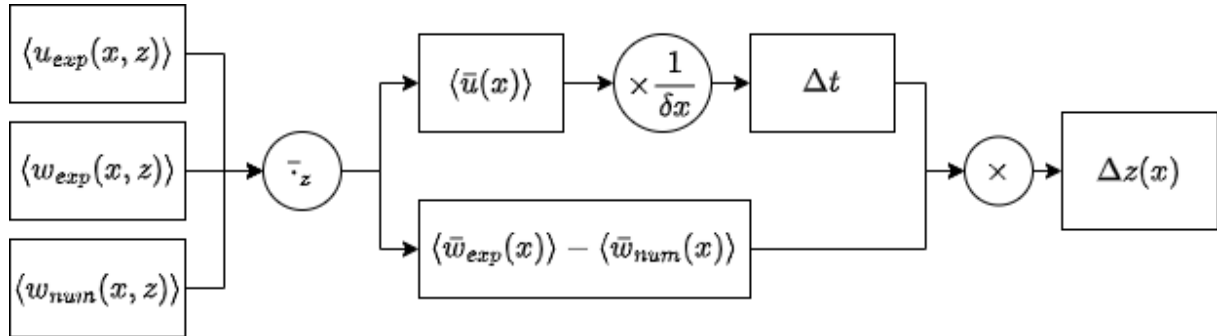


Figure D.1: Block diagram illustrating the method for applying a vertical position correction to the simulated wake relative to the experiments.

D.2. Vertical Shift Results

Figure D.2 shows the normalised time-averaged streamwise velocity profiles at eight streamwise locations behind the rotor. A vertical correction has been applied to each simulation profile (solid lines) based on the vertical velocity mismatch.

Previously, the vertical mismatch increased towards $x/D = 2.5$ due to growing vertical velocity differences. With the correction applied, the profiles are shifted downward, reducing the mismatch significantly. A small residual offset remains, likely due to averaging vertical velocity over the full vertical range instead of using pointwise values.

The correction also increases the separation between motion cases in the simulations. A clear distinction now appears between the high- and low-frequency cases, consistent with the experiments. The stationary case, however, remains very close to the high-frequency cases, showing minimal improvement.

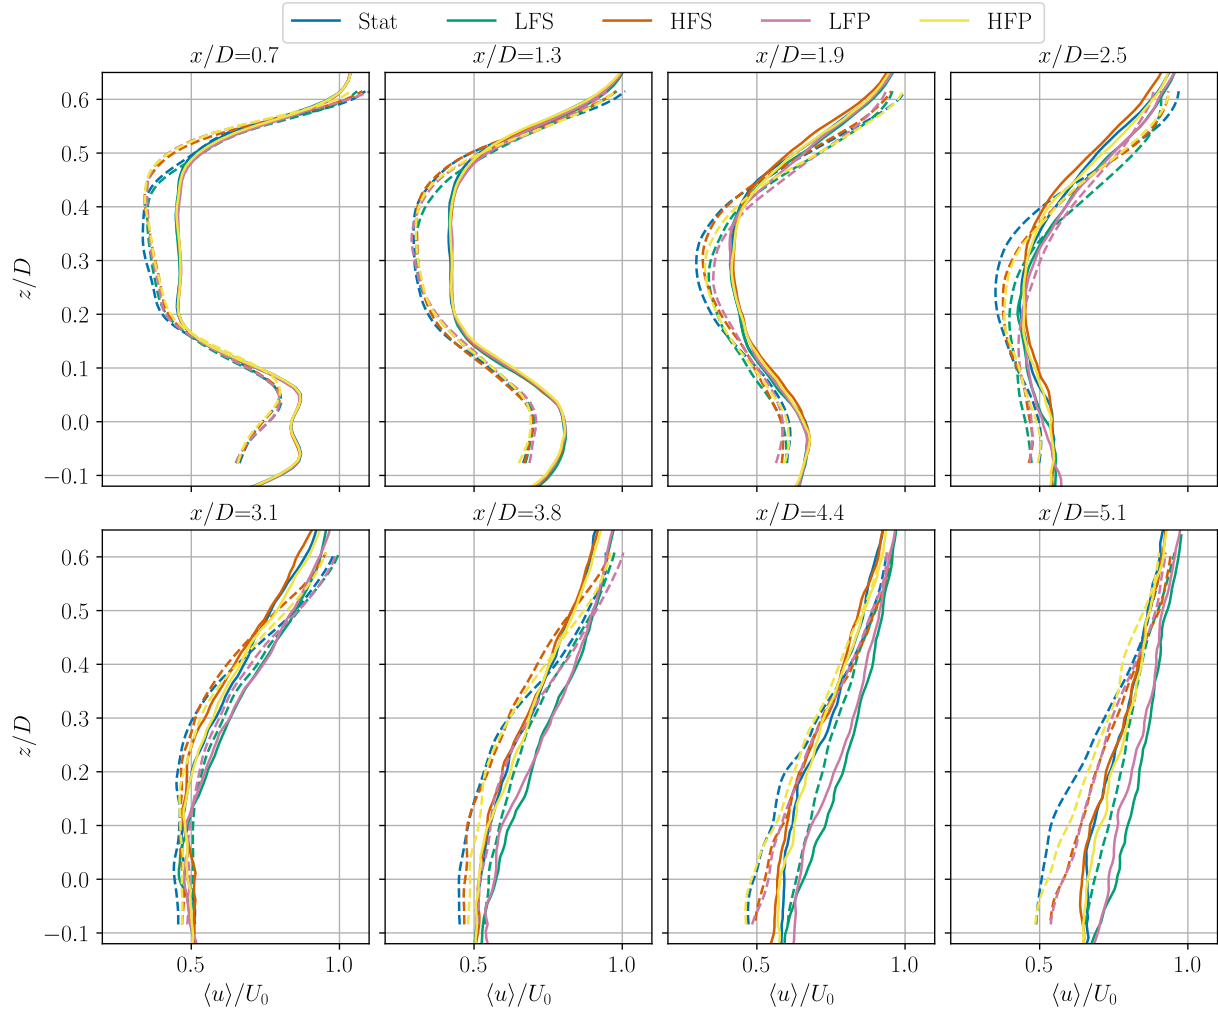


Figure D.2: Normalised time-averaged streamwise velocity along the vertical direction for simulations including a nacelle, after applying the vertical shift correction. Profiles are shown at $x/D = 0.7, 1.3, 1.9, 2.5, 3.1, 3.8, 4.4$ and 5.1 .