

Multidisciplinary Design Analysis and Optimisation for
Conceptual Design of Floating Offshore Wind Turbine
Support Structures

PhD Thesis

Katarzyna Patryniak

Naval Architecture, Ocean & Marine Engineering

University of Strathclyde, Glasgow

May 22, 2026

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed: Katarzyna Patryniak

Date: 16 May 2026

Published Work

The thesis consists of an introductory part and technical chapters based on a collection of published papers. The papers that form the basis of the thesis and a declaration of authorship are presented in the following.

List of Publications

The following six papers are considered in this thesis:

- P1** Patryniak K, Collu M, Coraddu A, Multidisciplinary design analysis and optimisation frameworks for floating offshore wind turbines: State of the art (2022). *Ocean Engineering* 251
- P2** Patryniak K, Collu M, Coraddu A, Multidisciplinary Design Analysis and Optimisation of Floating Offshore Wind Turbine Support Structures - Coupled Model Solution Strategies (2023). *Journal of Physics Conference Series* 2626
- P3** Patryniak K, Yeter B, Collu M, Coraddu A, Efficient design space exploration for spar Floating Offshore Wind Turbines: the trilemma of the ultimate, fatigue, and serviceability limit states (2024). *Journal of Physics Conference Series* 2767
- P4** Patryniak K, Collu M, Yeter B, Minisic S, Oneto L, Coraddu A, Hybrid frequency-time domain surrogate-augmented model of dynamics of floating offshore wind turbines (TBC). *Under review in Ocean Engineering*
- P5** Patryniak K, Collu M, Coraddu A, Rigid body dynamic response of a floating offshore wind turbine to waves: Identification of the instantaneous centre of rotation through analytical and numerical analyses (2023). *Renewable Energy* 218
- P6** Patryniak K, Collu M, Jonkman J, Hall M, Barter G, Zalkind D, Coraddu A, Investigation into Instantaneous Centre of Rotation for Enhanced Design of Floating Offshore Wind Turbines (2025). *Wind Energy Science* 10

Declaration of authorship

In all papers listed (**P1-P6**), K. Patryniak has taken the lead role in all aspects of the research and writing, assuming responsibility for conceptualisation, development of methodology, analysis, discussion, and writing. M. Collu and A. Coraddu provided valuable guidance and influence on the research ideas and methodology, offering insights that shaped the direction of the work and specific solutions. Additionally, in **P3** and **P4**, B. Yeter contributed by discussing and reviewing the work. In **P4**, S. Minisic, A. Coraddu and L. Oneto **performed the training of the surrogate models** using the concept and dataset generated by K. Patryniak, and followed by verification and implementation by K. Patryniak. In **P6**, J. Jonkman, M. Hall, G. Barter, and D. Zalkind contributed by discussing and reviewing the work, providing constructive feedback that enhanced the quality of the research. In all papers, all authors participated in revising the original manuscripts.

Abstract

Floating offshore wind turbines (FOWTs) enable wind energy harvesting in deeper waters, where wind resources are stronger and more consistent, supporting climate targets. However, their design is challenged by complex subsystem interactions, harsh metocean conditions, and competing performance requirements over long lifetimes. Current concept design typically relies on simplified models and sequential subsystem development, thereby limiting design scope and increasing the risk of overlooking globally optimal solutions. Despite recent advances, there remains a lack of a holistic, verified approach for integrated conceptual design optimisation.

This thesis aims to improve the efficiency and scope of support-structure conceptual optimisation by applying multidisciplinary design analysis and optimisation (MDAO) and a multi-domain, surrogate-supported numerical modelling framework. To support this aim, the optimisation problem is formulated in line with offshore standards, integrating serviceability, ultimate, and fatigue constraints, while remaining applicable to early-stage design. A novel resonance-avoidance constraint is introduced that is based on the probability and power content of excitation frequencies. A modular MDAO framework is developed to efficiently evaluate objectives and constraints. A hybrid frequency–time domain dynamic model is proposed, combining frequency-domain efficiency for motion prediction with time-domain accuracy for load and fatigue assessment. Key computational bottlenecks are mitigated by machine-learning models for hydrodynamic coefficients and mooring loads, applicable beyond a single project. Complementing the system-level optimisation, the thesis investigates fundamental FOWT motion characteristics via the instantaneous centre of rotation (ICR). A methodology

is developed to identify the ICR from numerical simulations. The results challenge a common assumption of a fixed centre of rotation and provide design insights.

The proposed methods are demonstrated through optimisation case studies, which illustrate efficiency gains, broader exploration of the design space, and new system-level insights. Collectively, this work establishes a foundation for efficient, integrated, and robust FOWT conceptual design, enhancing confidence in optimisation-based methods for next-generation FOWTs.

Contents

Abstract	iv
List of Figures	xi
List of Tables	xx
Acknowledgements	xxv
Abbreviations and Symbols	xxvii
1 Introduction	1
1.1 Motivation and background	1
1.2 Research aim and objectives	8
1.3 Main contributions to knowledge	11
1.4 Research design and approach	12
2 Literature review	19
2.1 Multidisciplinary design analysis and optimisation	20
2.1.1 Importance of MDAO	20
2.1.2 MDAO in FOWT design and other fields	23
2.1.3 MDAO methodology	26
2.1.4 An overview and comparison of available MDAO tools	28
2.2 FOWT design and optimisation	30
2.2.1 Design problem formulation	31
2.2.2 Unconventional metric: instantaneous centre of rotation	44

Contents

2.2.3	Design problem solution methods	49
2.3	FOWT numerical analysis approaches	55
2.3.1	Overview of modelling approaches	55
2.3.2	Time-domain models	58
2.3.3	Frequency-domain models	59
2.3.4	Simplified time-domain models	62
2.3.5	Extended frequency-domain models	63
2.3.6	Multifidelity models	71
2.3.7	Machine learning models	72
2.4	Summary and research gaps	75
3	Multidisciplinary design analysis optimisation framework	77
3.1	Introduction	77
3.2	MDAO framework	78
3.2.1	Coupled model	78
3.2.2	Fundamental disciplinary model	80
3.3	System solution strategies	86
3.3.1	Iterative solvers	87
3.3.2	Iterative solution example	89
3.4	Results and discussion	92
3.4.1	Influence of the introduction of the feedback loops	93
3.4.2	Influence of the choice of the iterative solver	96
3.5	Conclusion	99
3.5.1	Summary	99
3.5.2	Specific contributions	100
4	Optimisation problem re-formulation	101
4.1	Introduction	101
4.2	Design variables	104
4.3	Design objectives	109
4.4	Design constraints	110

Contents

4.4.1	SLS group	111
4.4.2	ULS group	121
4.4.3	FLS group	124
4.4.4	ALS group	127
4.5	Conclusion	127
4.5.1	Summary	127
4.5.2	Specific contributions	128
5	Unconventional metric: instantaneous centre of rotation	130
5.1	Introduction	130
5.2	Methodology	131
5.2.1	ICR derivation	132
5.2.2	Dynamic simulations	132
5.2.3	Environmental conditions	134
5.2.4	Floating system and coordinate system	135
5.2.5	ICR statistics	138
5.2.6	Analysis approach	140
5.3	ICR in waves	142
5.3.1	Regular waves	142
5.3.2	Irregular waves	153
5.4	ICR sensitivity to complex environmental conditions	156
5.4.1	Study setup	156
5.4.2	Results	157
5.5	ICR sensitivity to design variables	165
5.6	Responses sensitivity to ICR	169
5.7	Design application prime	170
5.8	Conclusion	172
5.8.1	Summary	172
5.8.2	Specific contributions	175

6	Hybrid-domain, surrogate-augmented model	176
6.1	Introduction and previous work	176
6.1.1	Efficient physics-based models	177
6.1.2	Data-driven models	181
6.1.3	Chapter outline	183
6.2	Methodology	183
6.2.1	FOWT system	184
6.2.2	Verification study setup	185
6.3	Physics-based model	187
6.3.1	Baseline coupled model of dynamics (C1)	187
6.3.2	Hybrid-domain model (C2)	200
6.4	Data-driven model	207
6.4.1	Potential hydrodynamics surrogate	208
6.4.2	Mooring surrogate	221
6.5	Case study	236
6.6	Conclusion	239
6.6.1	Summary	240
6.6.2	Specific contributions	242
7	Optimisation case studies	243
7.1	Introduction	243
7.2	Test 1: sequential versus integrated design	246
7.3	Test 2: the impact of fatigue constraints	257
7.4	Test 3: the impact of natural frequency constraints	262
7.5	Test 4: ICR-informed design	266
7.6	Conclusion	270
7.6.1	Summary	270
7.6.2	Specific contributions	271
8	Conclusion and future work	272
8.1	Context and problem	272

Contents

8.2	Aim and objectives	274
8.3	Novel contribution to knowledge	275
8.4	Future Work	278
Bibliography		282
A Aerodynamic damping		314
A.1	Precomputed look-up tables	315
A.2	Direct aerodynamic damping	315
A Nonlinear solvers algorithms and diagrams		318
A.1	Nonlinear Block Jacobi solver	318
A.2	Nonlinear Block Gauss-Seidel solver	320
A.3	Monolithic Newton's solver	320
B Nonlinear solvers examples		322
B.1	2D example	325
B.1.1	Jacobi iteration	325
B.1.2	Gauss-Seidel iteration	326
B.1.3	Hierarchical full-space Newton	326
C Flexible tower and modal analysis		328
D Tower structural checks		332
D.1	Shell (local) buckling	332
D.2	Column (global) buckling	335
D.3	Yielding check	337
E Verification numerical results for models C1 and C2		338
F Initial mooring surrogate dataset		340
G ICR time series in regular waves		342
H ICR in 1D and 2D coupled motion		344

Contents

I	Sensitivity of ICR to environmental loading parameters	345
J	Correlation between design variables and FOWT responses	347
K	Example of an IPOPT solution	349
K.0.1	Problem definition	349
K.0.2	Slack-variable formulation	349
K.0.3	Barrier problem and Lagrangian formulation	350
K.0.4	KKT conditions	350
K.0.5	Central path and convergence	351
K.0.6	Newton solution	352

List of Figures

1.1	Engineering design process phases [1].	3
1.2	Evolution of the knowledge, costs of change, and degrees of freedom during the design process. Reproduced from [2]	4
1.3	Modules interface in a fully coupled aero-hydro-servo-elastic model within OpenFAST. Reproduced from [3].	6
1.4	Thesis layout: the relationship between the research objectives, contributions, published work, and chapters.	12
1.5	High-level methodology.	13
2.1	Multilevel optimisation approach with a decomposition of the design space into two disjoint sets $\mathbf{x}_1$ and $\mathbf{x}_2$. Reproduced from [4].	26
2.2	An example of a feedback coupled scheme, illustrated following the XDSM method [5]. Disciplinary models output variables horizontally and receive inputs vertically.	27
2.3	Multidisciplinary Feasible MDAO architecture.	28
2.4	Example of mooring line profiles for the line shape parameters $x_M \in [-1, 2]$. Reproduced from [6].	33
2.5	Example frequency map for the IEA Wind 15 MW reference rotor adapted from [7]. Red curve - Kaimal wind spectrum, blue curves - JONSWAP wave spectrum for a range of sea states. Dark shading for the main range, light shading for the safety margins of ± 2 standard deviations. .	41

List of Figures

2.6	Cumulative Multi-Niching Genetic Algorithm (CMNGA) 2D design space exploration – multiple local maxima can be seen. Retrieved from [8] . . .	51
2.7	IEA 15 MW) controller regulation trajectory – variable speed variable pitch strategy. Adapted from [9]	70
3.1	Coupled model Design Structure Matrix - Optimisation loop.	78
3.2	Coupled model - Design Structure Matrix - Multidisciplinary Analysis (MDA) loop.	79
3.3	Panels convergence study. d -normalised panels density.	84
3.4	Comparison of the magnitudes of the surge and pitch RAO for the coupled and uncoupled models of dynamics. C1 - cycle 1, C2 - cycle 2. . . .	95
3.5	Surge RAO amplitude for different levels (sources) of damping: B_h - potential radiation damping, B_{aero} - aerodynamic damping, B_{add} - additional linear damping to calibrate with tank test, B_v - viscous damping. .	96
3.6	Coupling variables convergence history for different solvers.	98
4.1	Design parameters for the mooring line, floating platform and tower base section.	105
4.2	Geometrically plausible mooring lines range.	106
4.3	Geometrically plausible mooring lines.	106
4.4	Hydrostatics as a function of spar draft.	108
4.5	Serviceability and ultimate limit state representative wind speed, significant wave height and peak period selection based on conditional probabilities.	111
4.6	The process of deriving the natural frequency penalty for the first-order wave excitation region.	115
4.7	The process of deriving the natural frequency penalty for the wind excitation region.	116
4.8	1P frequency range penalty derivation process.	117

List of Figures

4.9	Comparison between the traditional (fixed region with $\pm 5\%$ margins [10]) and the new proposed resonance constraints for the rotor characteristic frequencies.	119
4.10	Wind, wave, 1P, and 3P range penalty functions with the reference system natural frequencies (dashed lines: surge, heave, pitch, and tower 1st fore-aft modes).	121
4.11	Fatigue limit state representative wind speed, significant wave height and peak period selection based on conditional probabilities.	124
4.12	Fatigue bin width sensitivity. 2 m s^{-1} bin marked in black.	125
5.1	Construction of the ICR using velocity vectors at two points on the floating body. Left: general case, misaligned nodes, middle: special case, aligned nodes, right: special case, pure translation.	133
5.2	FOWT design variables and coordinate system.	137
5.3	The procedure for investigating the ICR of FOWT, including an initial theoretical study (P5) and a more detailed study aimed at practical application (P6).	141
5.4	Wave energy spectra for 9 sea conditions (JONSWAP).	142
5.5	ICR time history in pitch natural frequency regular waves motion for 10 pitch periods.	143
5.6	ICR time history in pitch natural frequency regular waves motion for 1 pitch period, derived from Figure 5.5.	144
5.7	ICR space centre with probability distributions of the x and z coordinates. The coordinates of the ICR were z-score normalised (their distribution has a mean of 0 and a standard deviation of 1).	146
5.8	The expected value of ICR as a function of frequency for different wave heights.	147
5.9	Impact of nonlinear effects on $E(x_{ICR})$	148
5.10	Impact of nonlinear effects on $E(z_{ICR})$	149

List of Figures

5.11	Expected value of ICR as a function of frequency for different wave heights H , obtained for an unmoored system. The surge, heave, and pitch natural frequencies are marked with vertical lines for reference. . .	149
5.12	Comparison of the surge and heave displacement and velocity time histories for two different wave frequencies.	150
5.13	The time history and FFT plots of the pitch motion obtained with different dynamic models.	151
5.14	ICR time history during platform motion in a realistic environmental condition.	154
5.15	Histogram of ICR coordinates for a FOWT in a realistic environmental condition.	154
5.16	$E(x_{ICR})$ in irregular waves, for three frequencies and three wave heights. Box plots based on 30 random realisations.	155
5.17	$E(z_{ICR})$ in irregular waves, for three frequencies and three wave heights. Box plots based on 30 random realisations.	155
5.18	Effect of the regular wave period on the distribution of ICR coordinates (group A).	158
5.19	The effect of superimposing rotational velocity with (left) horizontal and (right) vertical translational velocities on the position of the ICR. . . .	159
5.20	Effect of the steady wind speed on the distribution of ICR coordinates (group C).	160
5.21	Effect of the wind shear on the distribution of ICR coordinates (group D).161	
5.22	Effect of the turbulence intensity on the distribution of ICR coordinates (group E).	162
5.23	Effect of the current speed on the distribution of ICR coordinates (group F).	163

List of Figures

5.24	Effect of the increased environmental loading complexity on the distribution of ICR coordinates. The regular/irregular wave cases only include regular/irregular wave loads; the steady wind case includes steady wind in addition to irregular waves; the complete environment includes irregular waves, turbulent wind, and current loads.	163
5.25	z_{ICR} range for different confidence intervals.	164
5.26	Skewness and kurtosis of z_{ICR} distribution for different environmental loading conditions.	165
5.27	ICR statistics for different one-at-a-time design variable perturbations. Original design values are marked with black outlines. Results are averaged over 10 environmental conditions.	168
5.28	Spearman correlations between the statistics of ICR and the design variables. Values closer to 1 and -1 indicate a stronger correlation. See Figure J.1 in Appendix J for correlations of all responses.	169
5.29	Spearman correlations between the statistics of ICR and the loads, obtained by varying the design features. Values closer to 1 and -1 indicate a stronger correlation.	170
5.30	ICR case study: design variables considered and their correlation with the ICR and responses statistics. Green/red arrows indicate a positive/negative correlation.	171
6.1	Dynamic models overview. PM–physical model; DDM–data-driven model; FD–frequency domain; TD–time domain.	184
6.2	Reference spar FOWT system illustration.	185
6.3	Wave rose for the NE8 site.	187
6.4	Force-displacement curves and length of chain on sea bed curves for the upwind and downwind mooring lines.	191
6.5	Verification of model C1 in operational cases. Mean loads. See Table 6.2 for reference.	196

List of Figures

6.6	Verification results in wave-only, wind-only, and operational cases. Standard deviation of motions and loads. See Table 6.2 for reference.	197
6.7	Mooring force-displacement curve obtained from the response to turbulent wind at different mean wind speeds and fixed turbulence intensity of 14 %.	198
6.8	The process behind model C2: from frequency-domain motion to time-domain loads.	201
6.9	Verification of model C2 in operational cases. Mean loads.	202
6.10	Verification of model C2 in operational cases. Standard deviation of the loads.	203
6.11	Verification of the responses - time series (left) and Fourier amplitudes (right) for in-plane platform displacements, nacelle horizontal acceleration, line 2 tension at the fairlead, and tower base bending moment. . .	204
6.12	Tension and stress cycles counting for the mooring line (fairlead). Case 3.1 (below rated conditions).	205
6.13	Tension and stress cycles counting for the mooring line (fairlead). Case 3.2 (rated conditions).	205
6.14	Tension and stress cycles counting for the mooring line (fairlead). Case 3.3 (above rated conditions).	205
6.15	Bending stress and stress cycles counting for the tower base section. Case 3.1 (below rated conditions).	206
6.16	Bending stress and stress cycles counting for the tower base section. Case 3.2 (rated conditions).	206
6.17	Bending stress and stress cycles counting for the tower base section. Case 3.3 (above rated conditions).	207
6.18	Physics-based and data-driven models integration. $\mathbf{x}$ - displacement vector.	207
6.19	Database generation process	208
6.20	BEM surrogate dataset input and output parameters distributions. . . .	211
6.21	BEM surrogate dataset output versus input parameters scatters.	212

List of Figures

6.22	BEM surrogate true–predicted scatter plots with absolute relative error histograms.	217
6.23	Diagnostic plots for $B_{h,33}$ prediction.	218
6.24	Diagnostic plots for $X_{im,3}$ prediction.	218
6.25	NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against waterline radius for different values of draft.	219
6.26	NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against draft for different values of waterline radius.	220
6.27	Comparison of the baseline FD model and the FD model with BEM surrogate in terms of the standard deviation of motions and loads. Performed for a reference design redefined with a reduced number of radii. See Table 6.9 for numeric results.	221
6.28	Mooring surrogate database generation process.	223
6.29	Mooring line parametrisation - parameters of the design of experiments and ML features.	224
6.30	Mooring surrogate dataset input and output parameters distributions.	227
6.31	Mooring surrogate dataset output versus input parameters scatters.	228
6.32	Mooring surrogate true–predicted scatter plots with absolute relative error histograms.	232
6.33	Diagnostic plots for standard deviation of F_z	233
6.34	True versus predicted mooring load component F_x	234
6.35	True versus predicted mooring load component F_z	235
6.36	Tension and stress cycles counting for the mooring line (fairlead), comparison between OpenFAST and the surrogate model. Case 3.2 (rated conditions).	235
6.37	Comparison of optimisation process with models C1, C2, and C3. The dashed curve indicates the original design.	238

List of Figures

7.1	Test 1 workflow: sequential system optimisation.	246
7.2	Test 1: Tower optimisation history. Black dashed line – initial value. . .	248
7.3	Test 1: Floating platform optimisation history. Waterline radius – r_4 . Black dashed line – initial value, blue dashed line – constraint limiting value.	248
7.4	Test 1: Mooring line optimisation history. Black dashed line – initial value, blue dashed line – constraint limiting value.	250
7.5	Test 1: Integrated support structure optimisation history. Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4	253
7.6	Test 1: Integrated support structure optimisation history initialised from the end point of the sequential optimisation. Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4	255
7.7	Test 1: Comparison between the initial and optimised FOWT designs. . .	256
7.8	Test 2: Support structure optimisation history with fatigue constraints (1st start point). Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4	259
7.9	Test 2: Support structure optimisation history with fatigue constraints (2nd start point). Black dashed line – initial value, blue dashed line – constraint limiting value.	260
7.10	Test 2: Comparison between the FOWT designs optimised with and without the fatigue constraints.	262
7.11	Test 3: Comparison of FOWT designs optimised with the natural fre- quency constraints locally from different starting points.	262
7.12	Test 3: Natural frequency penalty map for the FOWTs optimised locally from different starting points.	264
7.13	Design space visualisation: ICR and response statistics as functions of two selected design variables for the spar-type FOWT.	267
7.14	The responses of the modified spar FOWT design. Values normalised with the values of the original design as $1 + (x_{adjusted} - x_{original})/ x_{original} $	269

List of Figures

A.1	Aerodynamic force (thrust) and damping in surge direction curves for three values of mean pitch angle $\bar{\eta}_\theta$. The blade pitch (θ) and rotor speed (Ω) curves are shown for reference.	317
A.1	Nonlinear Block Jacobi solver. Reproduced from [11].	320
A.2	Nonlinear Block Gauss-Seidel solver. Reproduced from [11].	320
A.3	Full-space hierarchical Newton solver. Reproduced from [11].	320
D.1	Transition from yield to buckling strength.	334
D.2	EN 1993-1-1 Eurocode 3 buckling curves.	336
F.1	Initial mooring surrogate: true versus predicted value plots and absolute relative error histograms.	341
G.1	x-coordinate of ICR as excited by high-frequency (left) and low-frequency (right) waves of three different heights.	342
G.2	z-coordinate of ICR as excited by high-frequency (left) and low-frequency (right) waves of three different heights.	343
H.1	A comparison of pitch motion time history and FFT amplitudes in a 2 DOF coupled motion against a 1 DOF uncoupled motion.	344
I.1	Distribution of ICR coordinates in regular waves of periods 5 s–31 s (top) and 5 s–125 s (bottom).	346
I.2	Distributions of the horizontal velocity at points A (left) and B (right).	346
J.1	Correlation matrix between the design variables and the responses of interest for the spar design.	347
J.2	Correlation matrix between the design variables and the responses of interest for the semisubmersible design.	348

List of Tables

2.1	Comparison of the chosen elements of FOWT support structure MDAO frameworks - part 1. Acronyms explained in the <i>Abbreviations and Symbols</i> section.	21
2.2	Comparison of the chosen elements of FOWT support structure MDAO frameworks - part 2. Acronyms explained in the <i>Abbreviations and Symbols</i> section.	22
2.3	The MDAO tools	29
2.4	Objective functions proposed in FOWT design and optimisation studies.	39
2.5	Operational conditions constraint values proposed in FOWT design and optimisation studies.	45
3.1	Summary of modelling approaches within the analysis framework with reference to XDSM blocks.	81
3.2	Comparison of nonlinear solvers in terms of state updates and convergence behaviour.	87
3.3	Performance of the uncoupled and coupled models - the effect of cycle 1.	94
3.4	Performance of nonlinear iterative solvers. C1 - cycle 1, C2 - cycle 2. . .	97
4.1	Design variables and their bounds.	109
4.2	Cost coefficient values for structure and mooring components.	110

List of Tables

4.3	Environmental conditions. V_s - mean hub-height wind speed, TI - turbulence intensity, H_s - significant wave height, T_p - wave peak period, γ - wave shape parameter, θ_w - mean wave direction, V_c - current speed. ULS_1 and ULS_2 refer to conditions defined in Section 4.4.2.	112
4.4	Directionality correction factor calculation.	125
4.5	Summary of design formulation: variables, objectives, and constraints. .	128
5.1	Hydrodynamic model parameters.	133
5.2	Case study definition - wave conditions.	134
5.3	Environmental conditions binning. Wind shear fixed at 0.2 and turbulence intensity at 14 %.	135
5.4	OC3 Hywind spar FOWT original design characteristics. CG – centre of gravity; CB – centre of buoyancy.	136
5.5	NREL 5 MW reference rotor characteristics.	137
5.6	Spearman rank correlation calculation example.	139
5.7	Statistics of the ICR components based on the last 10 pitch periods, for the pitch natural frequency.	145
5.8	Impact of varied environmental loading on the ICR – conditions considered. γ – peak-shape factor.	156
5.9	The ICR of a FOWT in regular waves (group A).	158
5.10	The ICR of a FOWT in steady wind (group C).	160
5.11	The ICR of a FOWT in steady wind with shear (group D).	161
5.12	The ICR of a FOWT in turbulent wind (group E).	162
5.13	The ICR of a FOWT simulated in varying current velocities (group F). .	162
5.14	The ICR of a FOWT under environmental loading conditions of varying complexity. Each row represents one case of a group listed in Table 5.8. .	165
5.15	Original values and ranges of design variables. All values are in (m). . .	167
5.16	Summary statistics of responses based on the data used for the correlation analysis.	171
6.1	Mooring characteristics of the reference FOWT design.	184

List of Tables

6.2	Dynamic model verification cases.	186
6.3	Relevant S-N curves parameters.	194
6.4	Comparison of the stress and the resulting fatigue accumulation from a single environmental condition extrapolated to a full lifetime.	199
6.5	Ranges of the inputs of the dataset for potential coefficients surrogate training.	209
6.6	BEM surrogate dataset outputs (hydrodynamic coefficients).	210
6.7	Performance metrics used for surrogate model evaluation. y_i and $\hat{y}_i$ – true and predicted values of the output, $\bar{y}$ – mean of the true values, N – number of samples.	215
6.8	BEM coefficients surrogate aggregate error metrics grouped by mode. .	215
6.9	Relative errors due to incorporating the BEM surrogate model in the frequency domain baseline model (C1).	222
6.10	Mooring chain properties	224
6.11	Primary platform, mooring system, and environmental parameters (de- sign of experiments).	225
6.12	Mooring load surrogate model features.	226
6.13	Mooring load surrogate model targets.	227
6.14	Mooring model surrogate aggregate error metrics.	231
6.15	Average simulation time per design evaluation for models C1–C3 under 12 fatigue and 1 ultimate strength condition.	236
6.16	Initial and optimised designs comparison and verification. Values based on OpenFAST simulation results. Tension relates to the upwind line tension.	240
7.1	Test 1: Comparison between the initial FOWT design and the designs optimised in a sequential and integrated manner. Costs are normalised relative to the initial total cost of 3,452,500 USD.	256

List of Tables

7.2	Test 2: Comparison between the best FOWT designs from optimisations with and without fatigue constraints. Costs are normalised relative to the cost of 2,722,703 USD.	261
7.3	Test 3: Comparison of FOWT designs optimised with the natural frequency constraints locally from different starting points. Costs are normalised relative to the initial total cost of 3,452,500 USD.	263
7.4	Comparison of the original and the ICR-informed adjusted spar FOWT design responses. Body-fixed frame ICR values are in brackets.	268
7.5	Comparison of the original and the ICR-informed adjusted spar FOWT design responses - absolute and relative percentage differences.	268
7.6	The original and the ICR-informed adjusted spar FOWT designs features.	269
7.7	The original and the ICR-informed adjusted spar FOWT designs features - absolute and relative differences.	269
C.1	Tower sectional properties of the reference system evaluated at 11 stations.	330
C.2	Natural frequencies of the reference floating wind turbine system.	331
E.1	Verification results in operational cases for the baseline FD model (C1) against OpenFAST (OF). Mean of motions and loads.	338
E.2	Verification results in wave-only, wind-only, and operational cases for the baseline FD model (C1) against OpenFAST (OF). Standard deviation of motions and loads.	339
E.3	Verification results in operational cases for the FD+TD model (C2) against OpenFAST (OF). Mean of loads.	339
E.4	Verification results in wave-only, wind-only, and operational cases for the FD+TD hybrid model (C2). Standard deviation of loads.	339
F.1	Primary platform, mooring system and environmental parameters (design of experiments).	340
K.1	Example points on the central path for decreasing barrier parameter μ .	351

Acknowledgements

*Japanese **shīnsha** — profound and sincere gratitude, expressed with deep respect.*

I am grateful to the people I was lucky to have along the way — they are the primary reason I genuinely enjoyed the process and remained motivated and excited throughout.

To my family, friends, and colleagues old and new — for their belief in me, their patience, and for keeping me grounded. To Jason Jonkman, for our shared curiosity about the intricacies of floating wind turbine behaviour, for welcoming me to his team as a visiting researcher, and for his support in post-PhD pursuits. To my second supervisor, Andrea Coraddu, for the challenging and engaging discussions, and his careful attention to the quality of this work. To my supervisor, Maurizio Collu, whose guidance shaped this work and the researcher I have become. His intellectual rigour, willingness to challenge my thinking, and readiness to back ideas that others might have dismissed — and to give me the space to pursue them — steered this thesis to places I could not have reached alone. His patience and steady confidence made the difficult stretches of this journey navigable and even enjoyable. It is rare to find in a supervisor not only a trusted mentor who is genuinely invested in one's growth, but also someone whose company makes the hardest of problems a pleasure to work through. I could not have asked for better.

*Tra voi saprò dividere
Il tempo mio giocondo
Tutto è follia nel mondo
Ciò che non è piacer.
Godiam, fugace e rapido
È il gaudio dell'amore
È un fior che nasce e muore
Né più si può goder.
Godiamo, c'invita, c'invita un fervido
accento lusinghier.*

Abbreviations and Symbols

1P	Rotor frequency
3P	Blade passing frequency
A	Analytic
AEP	Annual energy production
BEM	Boundary element method
DOF	Degree of freedom
Emp	Empirical
FK	Froude-Krilov
FOWT	Floating offshore wind turbine
GD	Generalised displacements
ICR	Instantaneous centre of rotation
MB	Multibody
MCF	McCamy-Fuchs
Mat	Bill of material
M	Morison's equation
M+A	Material + anchors cost
MDAO	Multidisciplinary design analysis and optimisation
MDA	Multidisciplinary analysis
PBEM	Precomputed BEM
PF	Potential flow
PTD	Precomputed time-domain
QS	Quasi-static
RAO	Response amplitude operator
RV	Proportional to relative velocity
RNA	Rotor-nacelle assembly
ST	Strip theory
TLP	Tension leg platform
UBEM	Unsteady BEM

Chapter 1

Introduction

This chapter gives an overview of the motivation and background for this work and states the aim, objectives, and novel contributions to knowledge, as well as their relation to papers published by the author (or submitted and under review). It also provides a high-level overview of the research approach and maps the different components of the research process.

1.1 Motivation and background

Meeting climate targets while maintaining focus on the reliability and cost-effectiveness of energy has become a central concern for wind turbine engineers. Although land-based installations dominate the market, offshore wind capacity is expanding ten times faster than its onshore counterpart [12]. One of the primary limitations in the development of offshore wind farms is water depth; when depths exceed about 50-60 meters, bottom-fixed foundations may no longer be economically viable [13,14]. Floating offshore wind turbines (FOWTs) effectively address this challenge by placing a turbine on a floating platform with station-keeping capability provided by a mooring system.

FOWT opportunities and challenges Thanks to larger rotors and more consistent wind speeds, offshore wind offers greater capacity potential, which is crucial for increasing the share of renewable energy. The floating technology not only expands

the potential for offshore wind energy but also opens new possibilities for cost-effective operation and maintenance strategies. Passive or active adjustment of the floating wind farm layout for maximum power production and/or loads reduction [15] [16] and the tow to port strategy for major components replacement [17] are among the many developments enabled by FOWTs.

However, the lack of a rigid connection to the seabed and the significant motions of the moored structures introduce additional modelling, design and operational challenges. Both mean and dynamic hydro- and aerodynamic loads arising from waves and wind interact with the coupled motion of the system, leading to complex structural responses and failure mechanisms. Therefore, design of FOWTs requires structural integrity and performance assessment considering a diverse range of criteria and loading conditions.

Sequential design approaches In traditional floating support structure design for offshore wind turbines, the focus has often been on creating a highly stable and stiff platform, derived from the design principles of bottom-fixed offshore turbines [18]. While this approach ensures stability, it may result in a substructure that is heavier and more costly than necessary. In the early conceptual design phase, the design process typically relies on simple models and human intuition to generate various system concepts. These initial designs are then refined through subsequent stages, which involve more detailed simulations and analyses of particular subsystems.

Adding further complexity, FOWT development projects involve multiple companies and stakeholders with diverse expertise and interests [18]. Most notably, the design and manufacture of the floating platform and mooring system are handled by entities separate from the turbine and tower Original Equipment Manufacturers (OEMs) [19]. These two groups must then collaborate and iterate to develop a suitable structure and control system design — often in a manner that protects intellectual property (IP). Another developer typically receives the turbine-support structure assembly design as a given, plans the wind turbine array layout, grid connection, operations, maintenance, and decommissioning activities.




Figure 1.1: Engineering design process phases [1].

This sequential and manual approach is inherently slow and may overlook potential synergies or optimal solutions. If significant issues arise during later stages, designers may need to revert to the earlier design phase (like illustrated in Figure 1.1), leading to delays and inefficiencies [11]. Only a limited number of iterations through this process are typically possible. Therefore, the results are likely suboptimal.

Integrated design and optimisation This highlights the potential benefits of adopting more integrated and computationally advanced design methods that can address multiple system components simultaneously, thereby improving the overall efficiency and performance of the floating wind turbine systems [20,21]. By designing the entire system (the floating platform, mooring system, tower, power cable, rotor-nacelle assembly, and controllers) concurrently, it is possible to achieve an optimal balance between the performance and cost of all subsystems.

Further to that, casting the design problem as an optimisation problem enables the use of numerical methods to search for a globally optimum design by changing a possibly large number of variables (design features) across multiple subsystems, often subject to numerous constraints (design requirements) [11]. This approach allows for unbiased, systematic, and efficient exploration of the design space using appropriate numerical analysis models.

Although system-level optimisation has not traditionally been applied during the early design stages, it has been argued that the insights gained through optimisation at the conceptual design stage—when key decisions about the support structure topology are made—are crucial for the success of later phases. In later stages, design changes are more costly and difficult to implement (see Figure 1.2), making early optimisation




Figure 1.2: Evolution of the knowledge, costs of change, and degrees of freedom during the design process. Reproduced from [2]

decisions more impactful. As noted in [22] and [18], this is especially true for complex and innovative structures, such as FOWTs, for which there is limited prior empirical data available.

Design evaluation Unlike one-off analyses, where a less efficient model may suffice, multidimensional optimisation under multiple diverse constraints requires models that can provide quick evaluations to keep the process tractable. Specifically, efficient and accurate modelling is crucial to ensure that decisions made during optimisation yield reliable and effective designs. Poor efficiency can lead to impractically long optimisation times, while inadequate accuracy may result in suboptimal or unrealistic outcomes. Therefore, sufficiently advanced techno-economic models are essential to capture all design-driving effects while maintaining reasonable computational costs [23, 24]. The development of the analysis methods to inform the optimisation of FOWTs is inherently challenging due to several factors.

FOWT designs must be assessed over a typical operational lifespan of 20-30 years [22, 25], accounting for a wide range of environmental conditions and failure mechanisms. Risks associated with rare extreme environmental events (ultimate limit states), operational conditions (serviceability limit states), and progressive structural degradation over time due to recurring loads (fatigue limit states) must be considered. In the context of design optimisation, fatigue assessment poses a significant challenge, as it requires evaluating internal structural loads under a variety of environmental loading conditions (at least 20, according to DNV-RP-C203 [26]). The applicability and accuracy of efficient spectral-based fatigue assessment methods are limited by the broad-band loading

in offshore conditions, and time-domain approaches may be required, further increasing computational effort.

A full stochastic analysis of wind- and wave-induced loading on FOWTs over the lifetime using state-of-the-art, fully coupled nonlinear time-domain solvers with a considerable number of random realisations for each load case requires substantial computational resources, even for a single analysis [27, 28].

Perhaps most importantly, the challenge of FOWT optimisation lies in the complexity of the system and the broad range of specialised disciplines involved. The physical environment and dynamics of FOWTs are complex, with significant interactions between wind- and wave-induced responses [23]. The mean and dynamic fluctuating six degrees of freedom (DOFs) motion responses of FOWTs are strongly influenced by the magnitude of the aerodynamic loads, including both horizontal and vertical components [29, 30]. Platform displacements alter the position of the mooring system, which in turn causes shifts in the six DOF natural frequencies of the floating system [31–33]. Conversely, wave-induced platform motions significantly impact tower-top dynamics, affecting both the magnitude and variability of rotor thrust and torque [29, 30, 34]. Additionally, aerodynamic loads amplify coupling among the six rigid DOFs, with notable pitch-heave coupling emerging due to a positive mean pitch angle, altered hydrodynamics, and changes in mooring line tension [33, 35]. Large platform displacements can induce nonlinear behaviour (force-displacement relationship) of the mooring lines [29, 36] and alter the potential and viscous hydrodynamic load acting on the floating substructure [33, 37].

Given the strong interdependencies, analysing a FOWT requires contributions from multiple fields—including aerodynamics, hydrodynamics, structural dynamics, electrical engineering, control systems, and systems engineering, to name just a few. These disciplines are essential to fully model the floating system and capture interactions between various loads and the dynamic behaviour [24]. Many of these fields are also rapidly evolving, adding further complexity to developing and maintaining the coupled models.

Multidisciplinary design analysis and optimisation In this context, the multidisciplinary design analysis and optimisation (MDAO) methodology is highly beneficial, as it focuses on designing systems that involve multiple disciplines and subsystems [38]. MDAO is particularly useful for optimising complex, heterogeneous systems where interactions between disciplines are significant and cannot be overlooked, or where the synergy between subsystems can be leveraged [11]. This approach decomposes the model into smaller, highly specialised, interconnected analysis (disciplinary) blocks, each with its own set of input and output variables [39], while addressing the global system through variable coupling and an appropriate direct or iterative solution strategy [39].

To realise the fully coupled floating system, the disciplinary analysis tools within an MDAO framework must be appropriately interconnected. An exemplary way of interfacing the modules was demonstrated within the state-of-the-art dynamic analysis tool OpenFAST, as illustrated in Figure 1.3, where all disciplines, including the mooring, platform, tower, and RNA dynamics, as well as environmental forcing, form two-way (feedback) connections (dependencies), forming the so-called cyclic system.




Figure 1.3: Modules interface in a fully coupled aero-hydro-servo-elastic model within OpenFAST. Reproduced from [3].

The modular structure of an MDAO framework simplifies the development, modification, and maintenance of the complex model, while also facilitating collaboration among experts from different disciplines. It provides an ideal environment for computing total system derivatives, which are required for the efficient, gradient-based optimisation methods. Additionally, this model can be easily integrated with an optimisation driver, enabling flexible definitions of objectives, constraints, and variables. Any coupling variable passed between the components of the model can be treated as part of the optimisation problem formulation, offering the adaptability needed to address the specific challenges of diverse projects.

Disciplinary analyses The decomposition of a model into components enables multi-fidelity approaches, where disciplinary models can vary in efficiency and accuracy to best suit the specific problem or provide a choice of fidelity depending on the task at hand. While high-fidelity, well-validated coupled nonlinear time-domain models offer accurate results, they are often computationally expensive for iterative, holistic design optimisation. To mitigate this, approximate models can be applied. These can differ in several aspects: the domain (time domain, simplified time domain, frequency domain, or hybrid time-frequency domain), the number of degrees of freedom, DOFs (from fully flexible models to simplified rigid-body models), and the modelling paradigm (physics-based, data-driven, or physics-informed data-driven models).

These strategies can be combined to achieve the optimal trade-off between accuracy and efficiency across the full, coupled model. Specifically, while efficient physics-based frequency-domain approaches are typically accurate enough for global motion solutions, more computationally intensive time-domain models may be required to capture dynamic nonlinear motion-load effects in FOWT structures. Additionally, the most computationally expensive physical models can be replaced with precomputed look-up tables or data-driven (machine learning) models trained on mid- or high-fidelity data. These alternatives offer sufficient accuracy at a reduced computational cost, making them well-suited for highly iterative optimisation processes.

A novel design metric The strategies outlined above focus on improving the efficiency of the design optimisation, as guided by offshore standards and recommended practices [25]. While meeting these requirements is crucial in later design stages, the early conceptual design phase offers a unique opportunity for flexibility and creativity, and can be informed by metrics outside the standards and regulations. In particular, a fundamental understanding of the FOWT dynamic behaviour can be leveraged to develop novel design metrics that encapsulate the overall system performance without resolving complex load responses.

A key insight arose from the simple question of where a floating wind turbine pitches about—a curiosity that led to considering the instantaneous centre of rotation (ICR) as a potential design metric to guide the optimisation process. Rooted in basic principles of rigid body dynamics, the ICR offers an alternative way to understand the system’s behaviour. Identifying the instantaneous point where the floating system experiences minimum motion and strategically designing the FOWT considering ICR statistics can potentially improve FOWT performance, addressing multiple design factors previously formulated as separate design constraints.

1.2 Research aim and objectives

The aim of the thesis is to improve the efficiency and scope of FOWT support structure optimisation in the early design stages through the application of the MDAO methodology and a multi-fidelity, multi-domain, physics-based surrogate-supported model of dynamics. This research seeks to advance the understanding of the FOWT optimisation problem, improve its formulation, and refine solution methods, thereby enabling a more informed and systematic approach to FOWT design.

In the early design stages of FOWTs, there is a high degree of design freedom due to the multitude of engineering and operational decisions yet to be made (Figure 1.2). Traditional sequential design approaches often fall short in identifying the optimal configurations, as they rely on manual exploration of the design space, which is infeasible given its complexity. While formal design optimisation approaches offer

systematic searches for the best configurations, they are highly iterative. Moreover, a holistic design approach, as prescribed by the offshore standards, requires evaluating multiple objectives and constraints across various operational conditions, some of which demand complex modelling techniques, which can be computationally expensive and challenging to manage within optimisation frameworks.

To fully leverage the potential of optimisation methods and explore the extensive design space of multiple FOWT subsystems (in particular, the tower, floating platform, and moorings) across an extensive range of loading conditions, a practical approach to candidate designs evaluation is essential. This approach should balance efficiency and accuracy, incorporating prudent formulation and solution to the optimisation problem.

The first research objective of this thesis is related to the need to represent a complex FOWT system within an optimisation environment that is easy to manage and capable of handling and integrating multiple physical subsystems and analysis tools:

- O1** Apply a multidisciplinary design analysis and optimisation (MDAO) framework to the integrated optimisation of the FOWT system, coupling multiple disciplinary models within a unified and efficient optimisation workflow.

FOWTs operate in a complex offshore environment influenced by wind, wave, and current loads. As engineering systems, they consist of numerous subsystems, each with unique characteristics, behaviours, and modelling approaches. While technical aspects of system performance and safety are critical when designing FOWTs, other criteria—such as social and environmental impact and supply chain capability—are increasingly considered. A holistic evaluation of the entire system, encompassing both technical and non-technical disciplines, requires a variety of specialised analysis tools and approaches. This work integrates a range of reliable and proven analysis tools of varying fidelities into a coupled model, which is then verified and incorporated into a multidisciplinary optimisation framework. This framework ensures an efficient and transparent flow of information between modules, facilitates efficient coupled system iterative solution and total derivative computation for efficient gradient-based optimisation methods, and enables optimisation problem decomposition. Although this

integration currently focuses on techno-economic disciplines, the framework has been designed to be modular and extensible. As such, it provides a structural foundation for the future inclusion of additional non-economic considerations (e.g. following a PESTLE-type perspective [40]) even though their explicit integration lies beyond the scope of this thesis.

The second objective focuses on improving the accuracy to computational cost ratio for the coupled model of dynamics:

- O2** Develop, adapt, and verify disciplinary models suitable for conceptual design optimisation, achieving an appropriate balance between physical fidelity, validity range, and computational efficiency, in order to enable realistic and scalable FOWT design evaluations.

This includes improving the accuracy of load calculations through a hybrid frequency-time-domain model valid for a diverse range of FOWT designs under a wide range of environmental conditions, and improving solution time by incorporating surrogate models in place of computationally expensive parts of the model. The latter solution makes gradient computation throughout the model tractable, enabling a gradient-based optimisation approach that is typically significantly more efficient for optimising problems defined in a large design space with many independent design variables. An efficient model is also necessary to systematically study the design space and approach design optimisation in a more informed and holistic manner. Accordingly, the optimisation problem is formulated to reflect the decision-making processes characteristic of early-stage FOWT design, accounting for both natural frequency characteristics and forced responses, and considering ultimate and fatigue loads.

The last, third objective relates to a more efficient optimisation problem formulation:

- O3** Derive and test a novel design metric related to the dynamic behaviour of the FOWT system that simultaneously addresses multiple design requirements, simplifying the optimisation problem and supporting more informed conceptual design decisions

The motion of FOWTs is often viewed as a challenge, with significant effort devoted to minimising it. While minimising motion is important, this work takes a different approach by acknowledging that FOWTs are inherently dynamic systems. Instead of solely attempting to suppress motion, we explore how to leverage it to our advantage. By strategically modifying design features, we can influence the relative amplitudes, frequencies, and phases of surge, heave, and pitch motions to minimise dynamic displacements at critical locations on the structure. Contrary to the common assumption in the literature that a single point of minimum motion (i.e., the centre of rotation) exists, this research recognises that such a point can only be defined instantaneously, and its determination is nontrivial. This study develops a statistical approach to investigate the instantaneous centre of rotation (ICR) for FOWTs and explores how understanding its behaviour can inform the development of a new design metric. The central hypothesis is that this metric can integrate multiple concerns—typically addressed as separate design objectives or constraints—into a unified formulation, thereby simplifying the overall design optimisation process. This hypothesis is evaluated through the research presented in this work.

1.3 Main contributions to knowledge

Based on the research objectives defined in Section 1.2, several contributions have been made. While the specific contributions will be detailed at the end of Chapters 3 – 7, the key three contributions are listed below. Additionally, the links between the key contributions, published papers, objectives and the aim of the thesis are illustrated in Figure 1.4 (the list of published works is provided in the front matter of this thesis).

- C1** Development of a coupled frequency–time domain surrogate-augmented model of floating offshore wind turbine dynamics suitable for assessment of serviceability, ultimate, and fatigue limit states at the conceptual design stage.
- C2** Conceptualisation, dataset generation, integration, and verification of a widely applicable time-domain mooring surrogate model, agnostic to platform design and environmental conditions.




Figure 1.4: Thesis layout: the relationship between the research objectives, contributions, published work, and chapters.

C3 Investigation of the instantaneous centre of rotation of floating offshore wind turbines, addressing the limitations of common fixed-point assumptions, with development of a numerical model to analyse its statistics, design sensitivities, and design trade-offs.


1.4 Research design and approach

The high-level methodology is illustrated in the flow diagram of Figure 1.5 and detailed in the following.


Literature review Within this broad domain, the chapter reviews prior work and identifies key research gaps related to physics-based and data-driven coupled dynamic modelling, FOWT design methodologies, and optimisation problem formulation and solution strategies. Additionally, the problem of identifying the pitch centre of rotation is examined in the context of existing modelling assumptions and its role in FOWT design.

Literature review ----- Chapter 2


O1: FOWT MDAO framework ----- Chapter 3



O2: FOWT coupled model of dynamics ----- Chapters 4, 6



O3: Instantaneous Centre of Rotation of FOWT as a potential design metric ----- Chapter 5



Case studies ----- Chapter 7

Figure 1.5: High-level methodology.

MDAO framework To achieve the **aim** of improved efficiency and scope in the optimisation of FOWT support structures during the early design stages, MDAO methodology is adopted. This approach not only facilitates the construction of the optimisation framework (as formulated in **Objective 1**) but also supports the development of a coupled dynamic model required to evaluate design objectives and constraints (**Objective 2**). Rather than developing each disciplinary component from scratch, the methodology integrates existing open-source codes into a cohesive framework, focusing on establishing effective couplings between modules. This integration enhances the capabilities of each component through their interactions, creating a framework that offers greater utility than the sum of its individual parts.

The development of this framework progresses through several stages of increasing complexity, starting with a minimal frequency-domain coupled model and overall workflow, and followed by the incorporation of state-of-the-art components. At this stage, two significant gaps identified in the literature are addressed as follows.

Equilibrium convergence The first gap addresses the solution of the system’s static equilibrium in the rigid-body DOFs. At equilibrium, the product of the stiffness and displacement (offset) of the system equals the sum of the mean external forces applied to the floating platform. Two key static loads—the aerodynamic load and the mooring load—depend on the platform’s position. Specifically, the system’s pitch affects the rotor plane angle, influencing the aerodynamic load components, while the platform’s mean offset, primarily in the surge direction, affects the mean mooring loads.¹ As a result, both aerodynamic and mooring loads must be calculated at the platform’s mean position. However, the calculation of the mean position depends on these loads, creating a cyclic relationship between the mean loads and mean offsets. This requires an iterative solution approach. To the best of the author’s knowledge, most FOWT optimisation studies have neglected this coupling, solving for the platform’s mean position under mean external loads only once, without updating the loads based on the newly computed equilibrium position. This work introduces an iterative solution method to

¹This is specifically true for the floating platforms with catenary mooring systems.

resolve the coupling and determine the true static equilibrium.

Iterative solution strategy There are various strategies for solving coupled systems, with efficiency depending on the system’s structure (the balance between the computational effort of the components) and available computational resources (e.g., parallel processing). While the choice of solution strategy can influence overall solution time, the optimal approach has not been explored in the FOWT literature to date. This work benchmarks five iterative solvers, each differing in how coupling variables are managed.

Design problem formulation Any constrained optimisation problem can be formulated as the task of minimising a specific objective function with respect to a set of design variables, subject to a set of design constraints. The focus of this work is to improve the formulation of these constraints, enabling more realistic and less conservative designs from the earliest design stages. In this regard, the next stage of framework development involves enhancing its capabilities to tackle a more complex optimisation problem, formulated based on the offshore standards and recommended practices. While the early framework could analyse global platform motion and basic performance metrics such as maximum offset or nacelle acceleration, the improved model aims to explicitly evaluate structural loads. The structural constraints are assessed using the partial safety factor method, which simulates the load effects on the structure resulting from the applied loading processes.

In principle, FOWT structures must be designed to withstand loads under all temporary, operating, and damaged conditions to ensure safe operation, the safety of personnel (during maintenance activities), environmental protection, and adequate resistance to deterioration throughout the turbine’s design life. To meet these diverse requirements, FOWTs must be analysed in terms of a wide range of limit states—conditions within which the structure is expected to operate safely, and beyond which it is no longer able to meet the design criteria [41]. These limit states are classified into four main categories: ultimate, fatigue, accidental, and serviceability limit states, each encompassing various conditions that must be assessed throughout the design process.

This work focuses on selecting criteria that address the broadest set of requirements while still being evaluable with relatively cost-effective methods. These criteria are implemented in the optimisation framework, and their impact on the optimised designs is investigated.

This more detailed approach requires consideration of two key aspects:

- Formulation of the optimisation problem to better align with the recommendations of the standards and to address the loads directly,
- Development of computationally efficient methods for calculating the new constraints, suitable for the early design stages.

Hybrid domain model In particular, this work highlights limitations of purely frequency-domain approaches for evaluating loads in structural and mooring components. In several optimisation frameworks, quantities such as mooring-line tension are represented through response amplitude operators, from which statistical response measures are derived under the assumption of linear, stationary, Gaussian response processes. Design constraints may then be formulated using estimates such as the mean response plus a multiple of the response standard deviation. While this approach is computationally efficient and may be suitable for relatively narrow-banded response processes, such as loads dominated by wave-frequency excitation, under realistic environmental conditions, such statistical estimates may provide an incomplete representation of extreme and fatigue-relevant loads. To address this issue, this work transforms frequency-domain motions to the time domain and uses the resulting time series to compute dynamic loads more directly. The improvements offered by this hybrid approach, as well as its remaining limitations, are identified.

Implementation of surrogate models This work demonstrates two applications of data-driven models: one to replace a less accurate dynamic model component, and another to address an efficiency bottleneck. First, to overcome the limitations of the quasi-static mooring model in accurately predicting wide-banded mooring loads for fatigue analysis, a higher-fidelity time-domain code with a lumped-mass mooring model

is used to generate a dataset of motion and mooring load results. A machine learning algorithm is then applied to derive a functional relationship between the mooring configuration parameters and the fairlead motion time series (inputs), and the mooring tension time series (outputs). This approach enables the inclusion of mooring line fatigue damage as a constraint in the optimisation problem—something not previously achieved in FOWT MDAO studies. Secondly, a dataset of potential hydrodynamic coefficients is generated using a boundary element method for various floating platform configurations. This dataset is used to derive a relationship between the platform shape parameters and the hydrodynamic coefficients at different frequencies, enabling the replacement of the most computationally expensive element in the coupled model.

New natural frequencies constraint This work also proposes an alternative formulation for the natural frequencies constraint, based on probability and energy content. The FOWT system must be designed so that none of its critical natural frequencies overlap with the most energetic range of exciting frequencies, including wind and wave fluctuations, as well as rotor and blade passing frequencies (1P and 3P). In existing literature, such constraints are typically formulated in a stringent, conservative manner, often by imposing “hard” lower or upper (or both) limits on the system’s natural frequencies. Commonly, this means requiring the rigid body motion mode frequencies to fall between the peaks of the wind and wave spectra, and the tower frequencies to be in the stiff-stiff region, i.e., above the blade passing frequency. This work, however, acknowledges that load processes are characterised by site-specific probabilities of environmental conditions, along with the corresponding power spectral density. The proposed formulation aims to avoid the risk of resonance by considering both the probability and the expected severity of any potential resonance.

Instantaneous centre of rotation of FOWTs The final contribution of this work to the optimisation problem formulation is the derivation of an alternative design metric that simplifies the design optimisation process, as formulated in **Objective 3**. The dynamic behaviour of FOWTs involves complex interactions of multivariate loads from wind, waves, and currents, resulting in intricate motion characteristics. While

methods for analysing global motion responses are well-established, time- and location-dependent kinematics remain underexplored. A key aspect of this research is the investigation of the instantaneous centre of rotation (ICR) — the point of zero velocity at a specific time, corresponding to an instance of general plane motion. Although understanding and strategically positioning the ICR can reduce dynamic motion in critical structural locations, thereby enhancing the performance and structural robustness of FOWTs, it has not been explored in FOWT literature to date. This work develops a method to identify the ICR’s time series, statistics, and sensitivity to FOWT design features. It also investigates the application of the new knowledge of ICR in practical design scenarios, exploring whether this metric can be used to address multiple design constraints simultaneously.

Case studies The research concludes with a series of case studies that demonstrate the capabilities of the developed framework in enhancing the efficiency, scope, and realism of spar-type FOWT support structure optimisation (including the tower, floating platform, and moorings) during the early design stages.

This thesis focuses on the example of the OC3 Hywind spar platform [42] with the NREL 5MW reference wind turbine [43] – a configuration that was still industry-relevant at the outset of the study (2020) and remains extensively benchmarked and validated, making it well suited for methodological development.

Chapter 2

Literature review

This chapter is partly based on the published paper:

P1 Patryniak K, Collu M, Coraddu A, Multidisciplinary design analysis and optimisation frameworks for floating offshore wind turbines: State of the art (2022).
Ocean Engineering 251

Floating offshore wind turbine (FOWT) technology is still relatively new but is growing rapidly. Fifteen years after the installation of the first large-capacity floating wind turbine (Hywind, 2.3 MW), the most advanced FOWTs are now successfully completing the full-scale demonstration phase and achieving Technology Readiness Levels (TRL) 8 and 9 [44]. The largest floating wind farm to date, Hywind Tampen, with a capacity of 94.6 MW, has also been successfully installed and is now operational. Despite these advancements, significant design changes continue to emerge in newer designs, driven by lessons learned from earlier installations and ongoing research. These developments are further supported by advancements in numerical and experimental analysis techniques, a deeper understanding of the dynamic behaviour of FOWTs, and evolving design methodologies.

Between the first and most recent installations, research in FOWT design has undergone a significant transformation—from manual, brute-force cost optimisation of individual FOWT components to the integrated design optimisation of entire floating wind farms, incorporating a diverse range of constraints (i.e., multidisciplinary design

analysis and optimisation - MDAO). In this context, this chapter reviews the current state-of-the-art methods for multidisciplinary optimisation of FOWTs.

Section 2.1 introduces the MDAO methodology and reviews its applications in the FOWT and similar fields; Sections 2.2–2.3 provide a complementary review of early-stage design optimisation problem formulation, and of the state-of-the-art numerical analysis approaches in the context of multidisciplinary design optimisation.

This chapter frequently refers to Tables 2.1 and 2.2, which compare the design objectives, constraints, variables, and the analysis and optimisation methods used in the 17 FOWT MDAO studies. These studies were selected based on two criteria: consideration of the design variables across multiple subsystems and the inclusion of design constraints from multiple disciplines.

2.1 Multidisciplinary design analysis and optimisation

Multidisciplinary design optimisation (MDO) is a field of research that studies the application of numerical optimisation techniques to the design of engineering systems involving multiple disciplines or components [63]. In this thesis, the scope is focused on optimisation with respect to design variables spanning multiple subsystems (e.g., mooring system, floating platform, tower, control system), subject to constraints across various disciplines (e.g., hydrodynamics, aerodynamics, structural dynamics), necessitating a coupled analysis model.

2.1.1 Importance of MDAO

Introduction of MDAO strategies at the onset of the design process brings multiple advantages, including the ability to accurately represent the tight inter-dependencies between the subsystems, readily evaluate the system-level implications of technical decisions at the subsystem-level, and to encompass the stakeholders' requirements and lifetime considerations within a common design framework [18]. Through MDAO, the contributions of structural, hydrodynamic, aerodynamic, control, electrical, and power engineers, made concurrently, can be balanced to yield a truly optimal, safe system, not

Table 2.1: Comparison of the chosen elements of FOWT support structure MDAO frameworks - part 1. Acronyms explained in the *Abbreviations and Symbols* section.

First author Year	Sclavounos 2008	Fylling 2011	Myhr 2012	Hall 2013	Karimi 2017	Lemmer 2017	Dou2020 2020
Reference	[45]	[46]	[47]	[6]	[48]	[49]	[50]
Objectives and Constraints							
Economic	o	o	o	o	o	o	o
Performance (power)						o	x
Hydrodyn. (incl. motion)	o	x		o	o	o	x
Servo						x	x
Elastic (incl. fatigue)			x			o	
Mooring	o	x	o	x	x	x	x
Modal		x	x				x
Design Variables							
Platform	x	x	x	x	x	x	x
Mooring	x	x	x	x	x		x
Tower					x		
Cable		x					
Control						x	x
Optimisation Algorithms							
Algorithm	BF	SQP	B-A	GA	GA	PSr	SQP
Search goal	Global	Local	Local	Global	Global	Local	Local
Gradient-based		x					x
Derivatives		FD					A
Techno-economic Models							
Domain	FD	FD	FD/TD	FD	FD	FD/TD	FD
Aerodynamics	PTD	RV	UBEM	TD	TD	PBEM	PBEM
Wave excitation	PF	MCF+FK	M	PF	PF	PF	M
Hydrodyn. coeff.	PF	ST	M	PF	PF	PF	M
Viscous drag	-	-	M	M	M	Emp	M
Mean drift	-	Maruo1960	-	-	-	-	-
Slow-drift	-	-	-	-	-	-	-
Structures	R	R	FEM	R	R	MB	R
Mooring	QS	PTD	FEM	QS	QS	PTD	Al-Solihat2016
Cost	Mat	Mat	-	M+A	Mat	Mat	Mat

Table 2.2: Comparison of the chosen elements of FOWT support structure MDAO frameworks - part 2. Acronyms explained in the *Abbreviations and Symbols* section.

First author Year	Hegseth 2020	Hegseth 2020	Lemmer 2020	Hegseth 2021	Pollini 2021	Rohrer 2022	Rohrer 2023	Pollini 2023	Abbas 2024
Reference	[51]	[52]	[53]	[54]	[55]	[56]	[57]	[58]	[59]
Objectives and Constraints									
Economic	o	o	o	o	o	xo	o	o	o
Performance (power)									o
Hydrodyn. (incl. motion)	x	x	x	x	x	x	x	x	x
Servo									x
Elastic (incl. fatigue)	x	x	x	x			x	x	x
Mooring	x					x	x	x	
Modal	x	x		x	x			x	x
Design Variables									
Platform	x	x	x	x	x	x	x	x	x
Mooring	x				x	x	x	x	
Tower	x	x		x		x	x	x	x
Cable									
Control	x	x	x						x
Optimisation Algorithms									
Algorithm	S+S	S+S	BF	SNOPT	SQP	SNOPT	SNOPT	SQP	COBYLA
Search goal	Local	Local	Global	Local	Global	Local	Local	Local	Local
Gradient-based	x	x		x	x	x	x	x	
Derivatives	A			A	A	A	A	A	
Techno-economic Models									
Domain	FD	FD	FD	FD	FD	FD	FD	FD	FD/TD
Aerodynamics	PBEM	PBEM	PBEM	PBEM	PTD	PBEM	PBEM	PTD	BEM
Wave excitation	MCF	MCF	PF	MCF	M	MCF	MCF+FK	M	PF
Hydrodyn. coeff.	A	A	PF	A	A	A	A	A	PF
Viscous drag	M	M	M	M	M	M	M	M	Not spec.
Mean drift	-	-	PF	-	-	-	-	-	Not spec.
Slow-drift	-	-	[60]	-	-	-	-	-	Not spec.
Structures	GD	GD	MB	GD	R	GD	FEM	R	FEM
Mooring	[61]	-	QS	-	[62]	Spring	Spring	[62]	Not spec.
Cost	Mat	Mat	Mat	Mat	Mat	Mat	Mat	Mat	LCOE

dominated by a single design consideration [64]. Parallel computing and multi-fidelity approaches can be exploited [18], largely influencing the time of the solution.

To achieve the common objective of maximum lifetime cost of energy reduction, the specialists of all subsystems should work together. According to the findings of project FORCE [65], integrated design of wind turbines and their support structures is one of the most potent ways to save cost in offshore wind, capable of reducing the cost of electricity by at least 10%. However, currently, the rotor-nacelle assembly and the floating support structures are developed under separate contracts, with wind turbines identical to those used with bottom-fixed structures [18, 65, 66]. Before these commercial barriers are overcome, the focus is on the multidisciplinary design optimisation of the particular assemblies: support structure (discussed by this article) and RNA (for example [67, 68]), as a necessary step towards system-level, and eventually wind-farm-level, integrated design.

2.1.2 MDAO in FOWT design and other fields

The importance of MDAO can be well demonstrated by an example from aerospace industry where the introduction of composite materials triggered changes to the design process. Grossman et al. [69] showed that the integrated aerodynamic structural optimisation of a wing is superior to the traditional iterative design techniques due to the exploitation of the favourable interaction between the two disciplines. Namely, distributing the structural material such that large deformations did not reduce the aerodynamic performance (moving the centre of lift towards the root of a wing), hence relaxing the need to keep the deformations small. Reduced stiffness allowed for reducing the structural mass, increasing the overall performance of the wing in a way not attainable through a sequential design.

With a less apparent need for multidisciplinary optimisation when considering offshore floating structures, the MDAO concept was not broadly applied in that field. Only a few studies employed formal optimisation algorithms for the hull design (e.g., [70], [71]). Examples of multidisciplinary optimisation in this field are even more scarce. Perhaps the most advanced study of such type was conducted by Sugita and Suzuki [72]

who minimised the weight of a TLP subject to constraints stemming from different disciplines (mooring tension, platform offset and natural periods, structural stress).

The design of floating wind turbine support structures differs from the oil and gas platforms in many aspects. One of the most important distinctions is the fact that, while for O&G offshore structures the aerodynamic loads constitute a small fraction of the total load, FOWTs are designed to extract energy from the wind, and therefore are exposed to significant loads of wind and wind-generated waves. Despite this, initially, FOWT platforms heavily relied on the oil and gas industry legacy, with conservative designs.

The first study on multidisciplinary design optimisation of FOWTs was published by Scлавounos et al. [45], based on the thesis of Tracy [73]. The dynamic performance analysis was decomposed into three modules: mooring system, the floating structure, and the wind turbine, with the characteristics of each discipline assembled into one equation of motion based on a linear spring–mass–damper system in the frequency domain. Because the mooring loads may be nonlinear over large displacements, the platform steady state offset was computed by the hydrodynamic module and then used as a linearisation point by the mooring module, so that the offset and stiffness were computed iteratively.

After the world’s first FOWT, Hywind Demo (Karmøy), was deployed and proved to be a technically feasible concept, the research on support platforms gained more interest. Currently, the process of design and optimisation of these structures, and the urge to reflect the complexity of the system within the optimisation process, better resemble the practices seen in aerospace engineering than the traditional offshore sector.

One of the most sophisticated approaches to integrated design optimisation has been demonstrated by Hegseth et al. [52]. The authors optimised the Proportional Integral (PI) control system, tower, spar platform, and catenary mooring system in an integrated manner, by combining the structural and control state–space systems into a complete closed-loop aero-hydro-servo-elastic model. A modular approach was applied, with discipline-specific calculations performed by individual units of code, connected in a multidisciplinary network by feed-forward connections (with the exception of a feed-

back cycle for the calculation of viscous damping, as explained in [52]). By increasing the controller gains, a reduction in rotor speed variation was observed, however, this was achieved at the cost of increased fatigue damage. The cost and performance of the designs obtained through the integrated optimisation were also compared to those obtained without varying the controller gains, showing the superiority of the coupled approach over the simpler structural optimisation.

Ashuri et al. [4] performed a multidisciplinary design optimisation of the blades and tower of a bottom-fixed offshore wind turbine. The main challenges of the study lied in the fact that the analysis was performed in the time domain, multiple load cases, constraints, and variables were considered, and, on top of that, a gradient-based optimisation algorithm with finite-difference derivatives was utilised. To make the process attainable, a multilevel optimisation was utilised, where the design variables and constraints were decomposed into two levels (bi-level optimisation [74]). Up to the best knowledge of the author, this was the only distributed MDAO architecture FOWT MDAO optimisation study published up to date. As illustrated in Figure 2.1, each global optimisation iteration consisted of two optimisation processes at the lower levels, and the global iterations were repeated until convergence on the common objective was achieved. To further speed up the computation, 12 different design load cases were computed in parallel, each on a separate CPU.

A significant body of research on the MDAO of FOWTs emerged in the domain of control co-design, where the control system and the support structure are optimised concurrently. Bilgen et al. [75] developed a mixed-fidelity multidisciplinary control co-design optimisation framework where advanced controller's and blades structural parameters were adjusted to maximise the power coefficient while mitigating the loads on the structure. Abbas et al. [59] performed a series of optimisations, co-designing either the tower or floating platform with controllers. Du et al [76] co-optimised the platform and tuned mass dampers, reducing the platform mass, improving power generation, and reducing platform motions. Yu et al. [77] optimised the platform, the blade pitch controller, and the tuned liquid multi-column dampers (TLMCD) of a semisubmersible FOWT simultaneously. If tuned appropriately to the whole FOWT system dynamics,




Figure 2.1: Multilevel optimisation approach with a decomposition of the design space into two disjoint sets $\mathbf{x}_1$ and $\mathbf{x}_2$. Reproduced from [4].

TLMCDs provide additional damping hence lowering the platform pitch motion, which allows a larger blade pitch controller bandwidth and better dynamic response. In this case, the application of optimisation spanning multiple disciplines was therefore not only advantageous but necessary.

2.1.3 MDAO methodology

Multidisciplinary problems can be decomposed into a series of smaller blocks of computations (disciplinary solvers), each with its own set of input and output variables (“coupling variables”) [39]. The outputs of some of the blocks are passed as input to other components, creating either simple feed-forward connections or feedback connections, which create a coupled model (Figure 2.2).

The way the components of the coupled model are grouped and handled within an MDAO problem determines the hierarchy (or architecture) of the problem. MDAO architectures can be classified as monolithic, where a single optimisation problem is solved, or distributed, where the optimisation problem is decomposed into smaller,




Figure 2.2: An example of a feedback coupled scheme, illustrated following the XDSM method [5]. Disciplinary models output variables horizontally and receive inputs vertically.

independent subproblems [5].

Among the simplest monolithic MDAO architectures, and one that has been predominantly used in FOWT optimisation studies to date, is the multidisciplinary feasible (MDF) approach. In this method, disciplines are solved sequentially in a coordinated manner. When feedback connections are present, multiple iterations through the coupled model may be required for all disciplinary solvers to reach a consistent state. The MDF approach maintains separate disciplinary solutions, with only the final converged results being used to compute the design objectives and constraints handled by the optimisation driver.

The MDF approach is formulated in Equations 2.1 and illustrated in Figure 2.3 using the Extended Design Structure Matrix (XDSM) methodology [5]. In the equations, f_0 is the objective function, potentially dependent on multiple disciplines, x is the vector of design variables, y represents the coupling variables (outputs from disciplinary analyses), c_0 is a vector of design constraints involving multiple disciplines, and c_i is the vector of discipline-specific constraints. In the diagram, the green boxes denote the disciplinary components (analysis tools), while the parallelograms represent design variables (white) and coupling variables (gray) exchanged between components. Inputs are passed vertically, and outputs horizontally; variables above the diagonal feed forward, while those below feed backward.




Figure 2.3: Multidisciplinary Feasible MDAO architecture.

$$\begin{aligned}
 & \text{minimize } f_0(x, y(x, y)) \\
 & \text{with respect to } x \\
 & \text{subject to } c_0(x, y(x, y)) \geq 0 \\
 & c_i(x_0, x_i, y_i(x_0, x_i, y_{j \neq i})) \geq 0 \quad \text{for } i = 1, \dots, N.
 \end{aligned} \tag{2.1}$$

While the optimiser controls the design variables, the coupling variables are managed by a coupled model solver. The iterative solution of the coupled disciplinary model can be achieved using various nonlinear solvers, such as the Gauss-Seidel method, which solves one component at a time, or the Jacobi method, which solves all components and updates the coupling variables simultaneously (in parallel). More details on the coupled solution strategies will be provided in Chapter 3.

2.1.4 An overview and comparison of available MDAO tools

When constructing an MDAO framework, numerous challenges must be overcome, including the choice of the order of execution of the disciplinary analyses, the management of the flow of information between the components, the choice of the system solver (linear/nonlinear, monolithic/recursive iterators), and the integration with an optimisation algorithm. Efficient implementation of these aspects requires a very specific set of skills

and is time-consuming; the formalisation of system specification and workflow automation may take as much as 60–80 % of project time, according to the survey by Ciampa and Nagel [78]. In that respect, the open-source and commercial tools developed by third parties are very powerful, taking off a significant part of the coding burden. A brief review of the three chosen non-proprietary tools is given in the next section as summarised and compared in Table 2.3.

Table 2.3: The MDAO tools

Code	OpenMDAO	DAKOTA	DAFoam
Reference	[79]	[80]	[81]
Language	Python	C++	C++, Python
Gradient-based algorithms	SNOPT, IPOPT, SLSQP, NLPQLP, FSQP, PSQP, ParOpt, CONMIN	Conjugate gradient methods, SQP methods, Newton methods, MFD, Augmented Lagrangian method	pyOptSparse compatible solvers, SNOPT
Derivative-free algorithms	NSGA2, ALPSO	PS, Simplex, Greedy Search Heuristic, NOWPAC, EA, DIRECT, EGO	-

OpenMDAO [82] [79] is an open-source framework for efficient multidisciplinary optimisation developed by the MDO Lab (University of Michigan) in collaboration with NASA. One of its main advantages is the efficient data (input/output) passing between components through variables promotion or connect statements and a choice of system iterative solvers (Gauss-Seidel, Newton, and more). Highly valued is also its ability to efficiently calculate the total derivatives, either numerically (finite-difference or complex-step), or through analytic partial derivatives followed by the computation of total derivatives (through a direct or adjoint method). Hence, gradient-based algorithms can be applied to optimisation problems with a large number of variables and constraints efficiently. Being open-source, the code can be fully customised. Additionally, an interactive model structure visualisation tool is provided. An example of the application of OpenMDAO to FOWT optimisation can be found in [52].

DAFoam [81], developed at the MDO Lab, models multidisciplinary physics with

OpenFOAM – an open-source multiphysics software. The tool can deal with a large, constrained design space through the implementation of an efficient discrete adjoint method for total derivatives computation. The package also includes a geometry parametrisation module based on the Free-Form Deformation scheme (more details in Section 9), which makes DAfoam one of the most comprehensive tools available. However, the framework only supports optimisation algorithms available in the python library pyOptSparse, and the choice of physics solvers is limited to those implemented in OpenFOAM. An example of the application of DAfoam to engineering optimisation can be found in [83].

DAKOTA [80] is a multilevel parallel object-oriented framework for sensitivity and uncertainty analysis, design optimisation and calibration, developed by the Sandia National Laboratory with contributions from the community. This toolkit allows simulation codes to be interfaced with iterative mathematical and statistical methods through an interface developed by the user as a script in any language. This implies that DAKOTA can be connected to any simulation code, provided that it can be executed from a command line and performs its I/O through data files. A wide range of optimisation algorithms is available, including both gradient-based and derivative-free methods. Additionally, nested models and parallel computing can be enabled and easily managed. An example of the application of DAKOTA to engineering optimisation can be found in [84].

2.2 FOWT design and optimisation

Multidisciplinary design optimisation of FOWTs requires an appropriate formulation of the design problem in terms of defining the design goals and requirements, as well as an efficient method for solving the design problem. The literature on the approaches to these two aspects is reviewed in the following sections.

2.2.1 Design problem formulation

Many engineering design optimisation problems, including those of FOWT systems, can be classified as constrained multi-objective problems. The widely accepted mathematical formulation is given by [85]:

$$\min_{\mathbf{x}, \mathbf{p}} (f_1(\mathbf{x}, \mathbf{p}), \dots, f_k(\mathbf{x}, \mathbf{p})) \quad (2.2)$$

subject to the following constraints:

$$\mathbf{x} = [x_1, \dots, x_n]^T, \quad \mathbf{p} = [p_1, \dots, p_m]^T, \quad (2.3)$$

$$x_{i,L} \leq x_i \leq x_{i,U}, \quad i = 1, 2, \dots, n, \quad (2.4)$$

$$\mathbf{g}(\mathbf{x}, \mathbf{p}) \leq 0, \quad \mathbf{h}(\mathbf{x}, \mathbf{p}) = 0. \quad (2.5)$$

Where $\mathbf{x}$ is the vector of design variables, representing the system's parameters (e.g., structural dimensions, material properties, or geometric features), $\mathbf{p}$ is the vector of fixed design parameters, f_i is the i -th objective function to be minimised, $x_{i,L}$ and $x_{i,U}$ are the lower and upper bounds, respectively, for each design variable x_i , defining the feasible range of values. $\mathbf{g}$ and $\mathbf{h}$ are the inequality and equality constraints that must be satisfied for a feasible design.

The design variables, objectives, and constraints commonly formulated in FOWT MDAO studies are reviewed in the subsequent sections.

Design variables

To be able to analyse the performance of a large number of designs in an optimisation environment efficiently, their characteristics must be represented with the smallest set of variables possible. At the same time, sufficient design flexibility (or detail) must be ensured to represent the topologies adequately. Hence, a trade-off between the richness of the design space and the cost of the optimisation process is observed.

An MDAO application requires a common set of variables that can be manipulated and exchanged among different disciplines. Ideally, the geometric model should also

be smooth (shape modification should maintain a smooth geometry), provide local control, and short setup time. Unlike most of the geometric variables, the non-geometric features (such as material type, anchor type, etc.), and composition features (such as number of floaters, hull sections, mooring lines, etc.) often have discrete nature. Hence, the development of an efficient parametrisation scheme is a non-trivial task. See [86,87] for an extensive review of formal parametrisation schemes.

The design optimisation of FOWTs followed a less strict methodology for shape parametrisation. A very common approach is to only include the mooring line length and its orientation defined by two additional continuous parameters: the horizontal distance from the fairlead to anchor, and the angle of the lines with respect to a global coordinate system [88]. Some complexity can be added by allowing the diameter of the lines to vary along the line length [47] or by varying the material of each section of the line [89]. For the TLP support structures, the outer diameter of the tendons can be varied, as presented for example in [56] and [57].

A detailed mooring system parametrisation was presented in [51] and [55], where four variables of the mooring system were defined: line diameter, depth of the fairleads below the water surface, the total length of the line, and the horizontal distance between the fairlead and anchor. The inclusion of the mooring attachment point location in the set of design variables allowed the observation of the important platform-mooring system couplings. The fairlead depth below the water surface was shown to strongly influence the surge-pitch coupled motion, therefore influencing the pitch response of the entire system.

While most of the studies represented the mooring system by a set of geometric characteristic such as line length or anchor radius, Hall et al. [90] presented a unique approach, where the attachment of the lines to the hull was fixed, and the anchor positions were set to linearly vary with the parameter x_M ranging from 0 (i.e., anchor directly under the mooring attachment point) to 2 (i.e., the horizontal distance between the hull attachment and the anchor of twice the water depth). Taut configurations were represented by the negative values of the parameter x_M , as presented in Figure 2.4. For a given platform design and water depth, this parametrisation scheme allowed to




Figure 2.4: Example of mooring line profiles for the line shape parameters $x_M \in [-1, 2]$. Reproduced from [6].

model a continuous range of mooring configurations with just one design variable.

Floating platform parametrisation is greatly based on the examples of structures from the oil and gas industry. A few classical concepts exist, with spars, TLPs, and semisubmersibles being the most popular [91]. The geometry of each of the platform types can be uniquely defined by a different set of parameters. Therefore, the majority of the studies focuses on a chosen class, narrowing the range of free parameters down to a minimum.

The hull is usually modelled as either a single axisymmetric body or an array of these. For instance, Hegseth et al. [51] modelled the spar platform with 10 sections, each having a variable height, top diameter, and wall thickness, and Pollini et al. [55] modelled the spar with two diameters and the draft. The scantlings of the floating structures are usually considered only approximately by augmenting the wall thickness (for example in [6]) or material density (for example in [42]) to account for the additional mass of the internal steel members. However, when structural dynamics are included in a multidisciplinary optimisation problem, it is necessary to model these structures more accurately.

This was achieved, for example, by Hegseth et al. [51], who proposed to model the T-ring stiffeners inside each of the 10 sections of the floater with 5 additional variables: the thickness and length of the webs and flanges, and the distance between the stiffeners. To reduce the number of variables and the computational effort, the authors adopted a B-spline approach. By using 4 control points, smooth distribution of these parameters

along the depth of the floater was achieved, and the number of variables related to the stiffeners reduced from as many as 50 to just 20. A similar approach was adapted by Ojo et al. [92] who used a B-spline curve to represent the profile of the spar floating platform with 12 control points at predefined depths.

In their later study by Hegseth et al. [52], the authors noticed that the optimum wall thickness was mostly governed by the necessity to withstand the hydrostatic pressure. Hence, provided that the detailed scantlings design is not of interest, wall thickness can be expressed as a function of depth, with no need to be included in the set of variables.

Some floating platforms require a more detailed approach. For example, Pollini et al. [58] modelled the Tetra-spar floating platform with six variables defining the length and diameters of the central column and radial, diagonal, and lateral braces. Rohrer et al. [56] [57] parametrised the TLP platform with the draft of the main column, length of pontoons and their segments, and wall thicknesses of pontoon segments. Abbas et al. [59] parametrised a semisubmersible platform with the draft, platform radius, column diameter, and pontoon diameter. To dampen the heave, roll, and pitch motion in near-resonance conditions, heave plates are often used, particularly in semisubmersible designs. Gilloteaux and Bozonnet [93] allowed the use of multiple heave plates, with their spacing and radius as design variables. An interesting decision was to consider the ratio of the plate radius to the column radius, making this variable nondimensional. Lemmer et al. [49] additionally considered the thickness of the plates, showing that with thick plates the wave cancellation effect can be achieved, significantly improving the platform response in waves.

A unique approach was followed by the researchers at the University of Victoria [6, 48], where the design space covered a wide range of designs moving beyond the standard classification of the foundations. The eight-parameter scheme consisted of a single central cylinder (characterised by a variable draft, radius and taper near the waterline), accompanied by a variable number of outer columns (each characterised by a draft, radius, heave plate radius and distance to the central cylinder). Such an approach brings along multiple additional challenges, not seen when parametrising a single-class platform. The scheme must be able to represent both the existing design

topologies, as well as those not yet seen. The optimisation framework must be able to deal with the possible discontinuities in the complicated design space, where multiple local minima are possible. On top of that, the diverse types of geometries generated by the parametrisation scheme must be evaluated with one common analysis model with comparable accuracy.

Another approach to the task of covering all support structure configurations and beyond was presented by Hall et al. [90]. The basis function approach removes the consideration of the physical platform geometry and works with hydrodynamic coefficients instead. Hydrodynamic characteristics of each candidate design ($\mathbf{X}_0$) are represented as a linear combination of hydrodynamic characteristics of a set of basis platforms ($\mathbf{X}_i$), as given by Equation 2.6:

$$\mathbf{X}_0(\omega) = b_1\mathbf{X}_1(\omega) + b_2\mathbf{X}_2(\omega) + \dots + b_n\mathbf{X}_n(\omega) \quad (2.6)$$

where the hydrodynamic coefficients may or may not be functions of angular frequency, ω .

The optimisation problem is defined in such a way that the objectives are achieved by varying the coefficients of combination (or weights) b_i , which represent how much the final design resembles each of the basis designs. The final design is then reproduced by superposition of the basis geometries (no standard way of doing this exists). This approach avoids the requirement of expensive solution of the radiation and diffraction problems for every single design. Therefore, it offers a significant reduction in the optimisation time. However, ambiguity remains in the how these optimal combinations of basis functions should be translated to a physical geometry.

Note that the design of the fixed structural elements such as cross-bracings, tendon arms, ballast, fairleads, etc. was not discussed. Although essential for design integrity, such components are usually out of control of the optimisation algorithm and their parameters are chosen to fit the variable characteristics of each design. For an example of how these can be represented in a parametrisation scheme see [6].

The tower design is often represented as a truncated cone with variable base and

top diameters and a linear distribution of diameters in between, keeping the hub height fixed (e.g., [52]) or variable (e.g., [94]). Wall thickness at the tower base and top [58] or at multiple stations along the tower height [59] [57] are also commonly included in the design vector.

If the control system is subject to optimisation, additional design variables related to controllers are required. The most comprehensive parametrisation was demonstrated in control co-design study by Abbas et al. [59] where the design variables included the cut-off frequency of the low-pass filter, peak shaving percent, closed-loop bandwidth and damping ratio, parallel compensation gain and minimum sensitivity bandwidth.

Design objectives

This section focuses on the review and assessment of design objectives applied in the FOWT MDAO studies, starting with the economic metrics, followed by the performance metrics, as summarised in Tables 2.1 and 2.2.

The economic viability of a given wind turbine design can be assessed based on the LCOE, defined as the ratio between the total life cycle cost (sum of discounted CapEx, OpEx and DecEx) and total energy production (AEP) [95]. Since this factor reflects the cost of a unit of energy produced, lowering its value is beneficial for electricity consumers and improves the competitiveness of wind energy in the energy market. Lower LCOE can be achieved by either increasing the energy production or reducing the costs of financing, manufacturing, installation, O&M, and decommissioning.

LCOE was the objective applied in the study by Abbas et al. [59], where both the cost and performance parts of the formulation were accounted for Equation 2.7:

$$LCOE = \frac{(TCC + BOS) \times FCR + OpEx}{AEP} \quad (2.7)$$

The study used the NREL Cost and Scaling model [96] to obtain the capital expenditure (CapEx) as the sum of the turbine capital cost (TCC) and balance of system (BOS) multiplied by the fixed-charge rate (FCR). OpEx was assumed to be a constant rate per year, independent of the design. The inclusion of the annual energy production (AEP)

in the objective formulation allowed for quantifying the effect of the control strategy (peak shaving), which was subject to optimisation.

By far, most published studies use material weight and/or cost as the sole proxy for LCOE. Some positions also include the manufacturing costs [49, 51, 52] and the cost of anchors [6, 48]. The inclusion of the anchor cost is important when diverse FOWT designs are considered, as the gravity and driven anchors can be twice as expensive as the drag-embedded anchors [97]. An interesting approach was adapted by Tracy [73], who used the displacement volume as a cost driver, noticing that this single, easy to calculate factor accounts for both the mooring system static tension and the structural mass, which largely influences the capital cost. A similar approach is proposed in [98], with the structural mass approximated as the product of the displaced volume and an empirical factor. It is noteworthy that when the objective function represents capital cost or full LCOE, it is relatively straightforward to ensure continuity, smoothness, and differentiability, enabling the use of efficient gradient-based optimisation algorithms.

However, FOWTs must be both cost-effective and achieve the required response to ensure an efficient and safe energy extraction. For that reason, along with the cost, a series of performance objectives are sometimes considered (although not as often as purely economic metrics).

Sclavounos and Tracy [45] claimed the nacelle motion to be a critical performance measure. In this work, the standard deviation of the nacelle acceleration was expressed as an integral of the response spectrum (Equation 2.8):

$$\sigma_{\ddot{\zeta}}^2 = \int_0^\infty (RAO_{\ddot{\zeta}}(\omega))^2 S(\omega) d\omega \quad (2.8)$$

The authors argued that an excessive nacelle acceleration causes degradation of turbine performance and damage to the equipment. This metric was also considered important in [6, 48, 90], for slightly different reasons: large platform motions were thought to reduce the turbine lifetime, induce higher flapwise bending moments, and reduce energy production. Rohrer et al. [56] used nacelle acceleration as an objective to demonstrate the ability of the developed TLP optimisation framework to consider wave loading and

Response alongside the constraints related to wind loading and response. The popular belief that the nacelle acceleration is of uttermost importance in FOWT design was challenged by Nejad [99], who found there that there is no correlation between the tower-top acceleration and drivetrain responses. The study also concluded that the tower top acceleration does not affect the rotor torque (and so the power produced) in any significant way, as long as the pitch control system remains intact (similar conclusion can be seen for example in [100]). Finally, although the fatigue life was found to be the dominating parameter for the main bearings, it turned out not to be influenced by axial nacelle acceleration. These findings were applied by a relatively recent (2020) study by Hegseth et al. [51], who decided to discard this objective from their study.

Sclavounos and Tracy [45] demonstrated a multi-objective approach in which they aimed to minimise the nacelle acceleration, simultaneously minimising the mooring lines tension. A trade-off between these two objectives was exposed, showing that the minimum-tension designs (i.e., designs with catenary lines) were characterised by an increased nacelle acceleration. Myhr and Nygaard [47] aimed to minimise the force value in the mooring lines as well as their weight using a TD approach. Pillai et al. [89] performed an optimisation to investigate the trade-offs between the mooring system cumulative lifetime fatigue damage and its cost. Since mooring lines tension significantly influences both the dynamic behaviour of a wind turbine and the capital cost, the attention it receives in support structure optimisation studies is justifiable.

A different objective was considered in the study of Hegseth et al. [51], where the rotor speed standard deviation ($\sigma_{\dot{\phi}}$) was used as a proxy for power quality. For a given blade pitch angle, there is one optimum tip-speed ratio at which a turbine extracts the power from the wind most efficiently, i.e., at the highest coefficient of power (this can be seen in Figure 2.7). Therefore, for best performance, the ratio of the rotor speed to the wind speed should be kept at exactly this value (all else fixed). Large variations in the rotor speed at a constant wind speed result in suboptimal operation. Additionally, variations in the mechanical speed of the rotor (and hence the generator speed) necessitate the voltage and primary frequency control (to meet the grid connection requirements), which result in major power losses [101]. The need

to minimise the rotor speed variation was emphasised by numerous studies on control strategies, including [102].

Lemmer et al. [49] optimised a semisubmersible FOWT support structure with a unique objective of tower top displacement. The authors compared the results of two optimisations, one aiming to minimise the displacement on its own and the other aiming to minimise the displacement and LCOE, and demonstrated the trade-off between the two objectives. The first optimisation's optimal design had relatively high mass, resulting in low motions and low tower top displacement. The second, more comprehensive optimisation yielded a more cost-effective design but higher displacements.

The objectives considered in FOWT MDAO studies are summarised in Table 2.4.

Table 2.4: Objective functions proposed in FOWT design and optimisation studies.

Objective category	Objective function	Optimisation study	Ref.
Economic objectives	LCOE	Abbas et al., 2024	[59]
	Structural mass	Hall et al., 2013	[6]
	Structural and manufacturing cost	Lemmer et al., 2017	[49]
		Hegseth et al., 2020	[51]
	Displacement volume	Tracy, 2007	[73]
Motion & performance objectives	Nacelle acceleration	Sclavounos and Tracy, 2008	[45]
		Hall et al., 2013	[6]
	Rotor speed standard deviation	Hegseth et al., 2020	[51]
	Tower-top displacement	Lemmer et al., 2017	[49]
Mooring-related objectives	Mooring line tension	Sclavounos and Tracy, 2008	[45]
	Mooring tension and line weight	Myhr and Nygaard, 2012	[47]
	Mooring fatigue damage and cost	Pillai et al., 2019 (not MDAO)	[89]
Multi-objective	Nacelle acceleration and mooring tension	Sclavounos and Tracy, 2008	[45]
	Tower-top displacement and LCOE	Lemmer et al., 2017	[49]

Design constraints

The classification of the criteria as objectives or constraints is not consistent in the literature, with the same criterion acting as an objective in one study and as a constraint in another. For instance, the standard deviation of the nacelle acceleration, which was introduced as an objective in the previous subsection, was used as a constraint instead in [46, 50, 93]. Rotor speed variation presented as an objective in [51] was instead used as a constraint in the later work of these authors [52]. It was observed that control systems can reduce these fluctuations, however, it is done at a cost of an increased blade pitch actuator use. Therefore, the control effort was included in the set of constraints as well. The threshold values for these two constraints were based on scaled values characteristic for bottom-fixed wind turbines.

Likewise, the mooring line tension was defined as one of the design objectives in [45], but treated as a constraint in the remaining studies. Karimi et al. [48] applied the minimum line tension constraint to avoid excessive slack in the taut mooring system, while Pillai et al. [89], Hegseth et al. [51] and Brommundt et al. [88] set a maximum tension of the catenary line tension to avoid exceeding the breaking strength of the chains. The last two studies also constrain the loading on the drag-embedded anchors to be purely horizontal, to ensure the anchors remain effective, not being designed to take any vertical loading. The same goal was achieved in a different way by Pillai et al. [89], who set a limit on the number of line points in contact with the seabed, Dou et al. [50], who limited the length of the suspended line to be not more than 75 % of the total line length, and Pollini et al. [58], who constrained the vertical component of the tension at the fairlead to ensure that the mooring lines were never fully suspended. Rohrer et al. [57] constrained the pretension in the mooring lines to be positive, as ensured by the buoyancy exceeding the weight of the TLP system.

When the mooring system is subject to optimisation and both the taut and slack lines are considered, both the minimum and maximum line tension constraints are necessary to eliminate unfeasible configurations. This approach was followed in [73], [57], [46], and [58]. For instance, Tracy [73] formulated the maximum tension constraint as the sum of the mean and three times the standard deviation of the tension – representing

the dynamic component of the tension: $(T + 3\sigma_T)SF < T_{max}$. Similarly, the minimum line tension was formulated as $(T - 3\sigma_T)SF \geq 0$ – an easily applicable approach in frequency-domain frameworks.

A particularly useful constraint, which can be assessed early in the analysis process, is the avoidance of the resonance of the rigid body modes of oscillation with the environmental exciting loads. The full FOWT system must be designed in such a way that none of the important natural frequencies overlap with the most energetic range of exciting frequencies. Here, the wind spectrum, wave spectrum, rotor and blade passing frequencies (1P and 3P) are to be considered, as illustrated in Figure 2.5 for a 15 MW reference turbine [9].




Figure 2.5: Example frequency map for the IEA Wind 15 MW reference rotor adapted from [7]. Red curve - Kaimal wind spectrum, blue curves - JONSWAP wave spectrum for a range of sea states. Dark shading for the main range, light shading for the safety margins of ± 2 standard deviations.

The intention behind this constraint is to avoid excessive motions in resonance conditions. Often, this type of constraint is formulated more stringently by applying an additional 5-15% safety margin to account for the rotor speed variability and to ensure none of the system's important natural frequencies is near these ranges (as shown by the lightly-shaded regions in Figure 2.5). For instance, Myhr et al. [47] only discouraged the frequencies in the 1P and 3P ranges, while Ashuri et al. [94] considered wider near-resonant regions (i.e., $1.9P - 2.1P$ and $2.9P - 3.1P$). Pollini et al. [58] applied an inequality constraint to limit the surge and pitch motion and the tower frequencies,

with an additional margin of 15 %.

The excitation loads on a FOWT from wind and waves are characterised by the site-specific probability of occurrence of environmental conditions, along with the corresponding power spectral density of these loads. For instance, a site can be described by a probability distribution of mean wind speeds, where each mean wind speed corresponds to a specific power content (in general, higher wind speeds result in higher energy being transmitted to the FOWT system, subject to the control strategy). Similarly, FOWTs experience a range of sea states over their lifetime, each with its own probability and power content. Environmental conditions that are more likely to occur and/or carry higher power content pose a greater risk to FOWTs, as they can lead to more severe motions and structural loads. However, probabilistic formulation of the natural frequency constraints has not been demonstrated in FOWT literature to the best of the author’s knowledge.

In addition to the rigid motion modes, the resonance of flexible modes (e.g., tower fore-aft and side-to-side, blade flapwise and edgewise natural frequencies) with the exciting frequencies (such as wind, waves, 1P, and 3P) should be considered to avoid excessive structural loads. This has been addressed in several studies. For instance, Myhr et al. [13] utilised a nonlinear finite element method (FEM) to model 12 DOFs in the optimisation of a Tension Leg Buoy (TLB) platform. Dou et al. [50] and Hegseth et al. [51] considered the first tower fore-aft mode in their spar-type FOWT optimisation studies, solving the eigenvalue problem to obtain the natural frequencies. Both studies found that the optimisation shifted the tower mode above the 3P forcing region (i.e., to the stiff-stiff region). However, the methods for achieving this differed: the first study modified the floater and mooring system, while the second adjusted the tower design. Abbas et al. [59] constrained the tower’s first fore-aft frequency to be in the stiff-stiff range to avoid complex control behaviour. Pollini et al. [58] further explored the optimisation of both soft–stiff and stiff-stiff tower configurations, concluding that the specific target range of tower natural frequencies significantly impacts the optimal floating platform design and the overall cost of the system. This highlights the close coupling between floater and tower design and the requirement for their co-design.

In the case of spar platforms, it is also important to consider the Mathieu instability. This instability arises when the heave natural period is a multiple of half of the pitch natural period ($\frac{1}{2}, 1, \frac{3}{2}, 2, \dots$), leading to internal resonance and causing the pitch motions to become unstable [103]. Although the probability of this instability occurring is low, it remains a critical design consideration because, when it does occur, the consequences can be severe [104]. While this constraint can be design-driving, it has often been overlooked in the literature. One notable exception is the work of Hegseth et al. [51], who placed a constraint on the ratio between the heave and pitch natural periods. This constraint can be evaluated inexpensively using linear FD models, without directly modeling the nonlinear effects of the instability.

The structural constraints included the tower buckling [57,59], von Mises stress [59], as well as the fatigue damage in the steel structures of tower or floating platform [51–54,57–59] and mooring lines [46]. Note that while the inclusion of the tower fatigue became standard practice, only one FOWT MDAO study included fatigue of the mooring lines as a constraint. Accurate estimation of the mooring loads for fatigue analysis considering combined wind and wave loads is non-trivial and typically requires nonlinear dynamic models, which are not commonly applied in highly iterative optimisation studies. However, the dynamic motion of the floating system resulting from the oscillatory environmental loads can lead to increased fatigue damage accumulation in both steel support structures and mooring lines as compared to the bottom fixed counterparts [105,106]. Depending on the floating platform concept, the fatigue damage equivalent loads can be as much as 2.5-7 times higher than for land-based wind turbines [24]. As shown in [54], fatigue may be a design-driving constraint along the entire length of the tower and in the wave zone of the platform. Similarly, catenary mooring lines are prone to fatigue, being affected by the simultaneous presence of low and wave-frequency load and stress components [107]. In fact, the DNV-OS-J103 standard [108] states that the design life of a mooring line will often be shorter than the design life of the floating structure and the wind turbine, suggesting that the replacement of the mooring line during the design life of the floating system might be inevitable. Therefore, it seems worthwhile to include fatigue of both the tower and the mooring lines in

the design problem formulation.

Other constraints applied in FOWT support structure multidisciplinary optimisation studies include towing stability [58], static surge and pitch of the floating platform [55, 57, 58], the sum of the static and dynamic pitch motion evaluated in the maximum thrust condition [6, 48, 93], the standard deviation of the free surface elevation to avoid slamming and green water [73, 93], maximum platform motion to avoid breaking the power cable [88] or maximum cable tension [46], centre of gravity relative to the centre of buoyancy to preserve the gravity-stabilised type of the support structure [58], rotor over-speed [59]

Table 2.5 lists the constraint values assumed in FOWT optimisation studies. It is important to note that some level of ambiguity arises in setting these limits. Most published FOWT optimisation studies constrain motion parameters (displacements and accelerations) rather than directly constraining the resulting structural loads. For example, studies limit the maximum surge offset to protect the power export cable, rather than directly limiting the cable’s structural load. Setting these motion-based limits arbitrarily can lead to overly conservative designs. In contrast, directly addressing structural loads offers the potential to reduce system costs while maintaining safety and performance.

2.2.2 Unconventional metric: instantaneous centre of rotation

The dynamic behaviour of FOWTs presents significant complexity due to the diverse range of coupled loads they encounter in an offshore environment. The distinctly distributed inertial, hydrostatic, mooring, and aero- and hydrodynamic loads are influenced by varying combined wind, wave, and current impacts. It is therefore of interest to formulate a compact metric that encapsulates the system’s overall dynamic behaviour. The instantaneous centre of rotation (ICR) could potentially serve as such a metric, as discussed in the following.

Table 2.5: Operational conditions constraint values proposed in FOWT design and optimisation studies.

Constraint	Units	Limit	Optimisation study	Reference
Max static pitch angle	°	10	Hall et al., 2013	[6]
		10	Karimi et al., 2017	[48]
		5	DNV-RP-0286	
Max static + dynamic pitch angle	°	10	Tracy, 2007	[73]
		10	Hall et al., 2013	[6]
		10	Karimi et al., 2017	[48]
		9	Fylling et al., 2011	[46]
		15	Hegseth et al., 2021	[54]
		10	Dou et al., 2020	[50]
		10	Gilloteaux et al., 2014	[93]
		5	COREWIND	[109]
		7	NORCE	[25]
		10	DNV-RP-0286	[10]
Max dynamic pitch angle	°	1	COREWIND	[109]
Max RNA acceleration	m s^{-2}	2.6	Fylling et al., 2011	[46]
		2.0	Dou et al., 2020	[50]
		1.0	Karimi et al., 2017	[48]
		5.0	Gilloteaux et al., 2014	[93]
		1.85	COREWIND	[109]
		2.94	NORCE	[25]
		2.94	DNV-RP-0286	[10]
Max horizontal offset	% of depth	15	Fylling et al., 2011	[46]
		10	Hegseth et al., 2020	[51]
		15	Dou et al., 2020	[50]
		15	COREWIND	[109]
Min line tension (TLP)	kN	0	Tracy, 2007	[73]
		0	Hall et al., 2013	[6]
		1000	Myhr et al., 2012	[47]
		728	Fylling et al., 2011	[46]
		0	Karimi et al., 2017	[48]

The concept and misconceptions

When a FOWT undergoes general plane motion, one can identify a point, known as the instantaneous centre of rotation (ICR), where the velocity is zero at a given mo-

ment [110]. This point is not necessarily located within the physical boundaries of the FOWT (and may even lie outside the finite physical domain; for example, pure translational motion corresponds to an ICR at infinity) and is highly dynamic, shifting in response to changing environmental loads and the resulting platform motions. Accurate identification of the ICR and its statistics can offer valuable insights for optimising motion reduction strategies and enhancing the FOWT’s performance and structural robustness.¹

By strategically modifying design elements to influence the position of the ICR, dynamic motions at critical locations can be minimised. For example, aligning the fairleads with the ICR could limit the dynamic part of the motion, which does not contribute to useful restoring but does contribute to fatigue damage of the mooring lines. Similarly, positioning the ICR near the rotor-nacelle assembly (RNA) could decrease the rotor and nacelle motion, reducing aerodynamic load fluctuations and mitigating fatigue throughout the system, as well as reducing human exposure to motion, improving the safety of the onsite maintenance operations. Lastly, the center of rotation (CR) does not always align with the centroid of the waterplane area (centre of flotation) [111], except in cases of small quasi-static rotations with no dynamic forces on the floating body. As a result, points at different distances from the CR experience different vertical displacement amplitudes, meaning that — even for a geometrically symmetrical structure — opposite sides of the platform undergo unequal vertical motions during pitching [112]. This spatial variation in response may inform the design of damping plates for semisubmersible platforms, allowing optimisation to leverage the resulting load imbalances.

Despite the potential benefits, the ICR of FOWTs has received very little consideration in the floating offshore wind community. The majority of the literature focuses on static or quasi-static scenarios, with the assumption of small motions around an equilibrium position and a fixed centre of rotation. Wen et al. [113] studied the effect

¹Note that identification of ICR is not strictly required to compute the responses with fully-coupled aero-hydro-servo-elastic nonlinear time domain dynamic solvers, in which case the physics of the problem (distributed loads) drive the ICR, and not vice versa. The knowledge of ICR is useful for lower-fidelity numerical models that make assumptions about the centre of rotation (CR), as well as to understand the FOWT coupled motion patterns.

of the combined surge and pitch motion of a spar-supported WT on the effective rotor velocities, assuming the pitch centre to be constant in space and time; the exact position has not been reported. Tran et al. [114] and Liu et al. [29] studied the coupled dynamics of FOWT using Computational Fluid Dynamics approaches. Tran et al. [114] assumed a fixed centre of rotation at 90 m below the hub centre, noticing, however, that for different foundation concepts, this assumption might not hold. Liu et al. [29] (and Johlas et al. [115]) assumed the system rotated about an axis through the centre of mass. Kelberlau [116], on the other hand, assumed the CR of a semisubmersible floater to be located at the midpoint between the centre of gravity and the centre of buoyancy.

Although these assumptions allowed significant progress in understanding the responses of FOWT and led to the development of useful analysis and design tools, the field still lacks an appreciation of the unsteady behaviour of these structures. The literature often refers to *the centre of rotation*: a single point (or axis) about which the system is assumed to rotate. While such a point can indeed be defined at a particular instant in time, for example, at an instant when the platform is in equilibrium, this gives no insight into the platform motion patterns during the motion. To the best of the author’s knowledge, no methods to study and describe the highly complex time-dependent rotational motion of the FOWTs have been developed up to date.

ICR in other fields

Early studies on the centre of rotation of floating structures were conducted by Steward et al. [117] and Standing et al. [118], focusing on barge vessels. Both authors noted significant oscillations in the centre of rotation, with Standing observing that the point “tends to wander wildly” in irregular waves. This work challenged the assumption that the centre of rotation is aligned with the centre of gravity, and highlighted the importance of hydrodynamic loads. Haslum et al. [104] further examined the centre of rotation in a spar platform undergoing surge-pitch motion, revealing its sensitivity to oscillation frequency and fairlead location. Souza et al. [119] introduced an experimental method to determine the ICR using velocity vectors from two probes on the floating body. Fernandes et al. [111] and Costa et al. [120] built on the previous work by

analysing a floating production storage and offloading (FPSO) vessel in regular waves, showing that the ICR locus (a curve formed by all ICR points) is a straight line that does not align with the vessel's symmetry and varies with wave frequency. At very low frequencies, the ICR exhibited a behaviour akin to pure translation, approaching an infinite distance above the body. Costa et al [121] found that the ICR coordinates of a moored vessel in regular head waves, free to move in surge, heave, and pitch, closely follow a Cauchy distribution. In subsequent work [122], the authors expanded their investigation to oblique seas and six DOFs, further reaffirming the ICR's dependence on external loading frequency.

ICR in FOWT field

Despite its relevance, the ICR has not been a focal point of FOWT studies. A few examples include the work by Eliassen et al. [123], which was the first attempt to investigate the centre of rotation of a FOWT without presupposing its location, and the study by Kaptan et al. [124], which examined the influence of wave frequency on the centre of rotation for two different floating concepts. Eliassen et al. [123] calculated the z -coordinate of the centre of rotation based on the mass, added mass, and stiffness properties through the ratio of two eigenmodes. Kaptan et al. [124] identified a point where the total horizontal motion with contributions from the surge and pitch motions is zero. Following the observation that the z coordinate of such a point satisfies $\eta_1 + z\eta_5 = 0$,² the CR was obtained based on the ratio of the complex transfer functions between the surge and pitch motions and the wave amplitude. Almost no dependence on the wave frequency for the spar platform and a significant, non-monotonic dependence for the semisubmersible platforms were indicated. Similarly, Lemmer et al. [53] considered the centre of rotation to be the point where the surge and pitch response amplitude operators combine to produce near-zero motion. The authors utilised taut mooring semisubmersible designs to ensure that the system rotates about the hub, thereby achieving preferable dynamic response characteristics, improved power quality, and

²That is, a combination of pure translation η_1 and the translation due to the rotation η_5 of the system's origin around the centre of rotation, should be equal to zero

reduced tower base loads. While the location of the centre of rotation was determined through frequency domain analysis, its instantaneous nature was not taken into account.

2.2.3 Design problem solution methods

The design process for floating wind turbines (FOWTs) is typically constrained by strict budgets and timelines, necessitating efficient design approaches (Korch, 2024). As a result, a significant volume of research has focused on improving FOWT design methodologies. This section investigates the FOWT design approaches presented in the literature, and then zooms in on optimisation approaches, focusing on the optimisation algorithms and derivatives computation.

Design approaches

The literature reveals a shift from sequential design processes, through simple parametric design searches and optimisation of individual components, to more holistic parametric design studies formulated as optimisation problems, and integrated optimisation approaches. As the size and complexity of the problem increase (e.g., more design variables and constraints), parametric studies are less likely to identify globally optimal solutions, making advanced optimisation algorithms increasingly useful. Despite this, parametric sizing studies are still commonly performed, either within a limited design space (narrowly bounded design variables) or with a strong emphasis on the efficiency of analysis models.

For example, Souza et al. [125] conducted a parametric design and structural analysis of three 20 MW spar-type FOWTs, adapting an initial floating platform design previously optimised for a different rotor, along with a previously optimised land-based tower. The floating platform designs were modified to satisfy requirements for draft, floatability, static pitch angle, metacentric height, and system natural frequencies. The tower design aimed to achieve natural frequencies in the stiff-stiff region (with frequencies larger than the blade passing frequency), a design approach found feasible in other FOWT optimisation studies. A time-domain, nonlinear aero-hydro-servo-elastic model was used to assess sectional fatigue damage, extreme motions, and axial stresses.

Pan et al. (2024) proposed an integrated design methodology for shared mooring systems for FOWTs, which begins with an initial parametric study of individual mooring systems under constant thrust conditions, followed by an array-level investigation of the impact of shared mooring designs on the performance of the floating system under different environmental conditions. Although both studies were framed as optimisation problems, with the design variables, objectives and constraints defined, neither applies a formal optimisation algorithm, instead relying on a design space search and selection approach to identify the most cost-effective configurations that satisfy various design constraints.

The recent (2024) report from IEA Task 49 [25] proposes a two-stage floating wind farm design process: i) design of the individual components (e.g., mooring lines, dynamic cables), and ii) integration of these components into a wind farm layout, followed by adjustments for array-level considerations. This top-down approach begins with the sequential design of subsystems and their integration, followed by an iterative process of refining both high- and low-level details of the wind farm design. During the component selection phase, system-level constraints are already taken into account—component designs are considered “forward-looking” in terms of their expected performance when integrated into an array situation. Importantly, the authors propose a common language for discussing the integrated design of FOWTs, using the nomenclature of an optimisation problem, including design variables, objectives, and constraints. The method for solving the design problem is not specified, as this problem formulation is suitable for one-off designs, parametric studies, and formal design optimisations.

If the design problem is to be solved by an optimisation algorithm, the selection of such algorithm is governed by the consideration of the type of optimisation problem at hand, the search goal, and the search method.

Optimisation algorithms

Like the majority of engineering problems, FOWT optimisation problems can be categorised as linearly or nonlinearly constrained problems, therefore requiring optimisation algorithms that can handle constraints. In some studies, the constraints are included




Figure 2.6: Cumulative Multi-Niching Genetic Algorithm (CMNGA) 2D design space exploration – multiple local maxima can be seen. Retrieved from [8]

by introducing penalty functions³, so that the use of unconstrained minimisation techniques can also be made (e.g., [47]).

The search goal defines the goal of the optimisation – whether the global or local optimum/optima are sought. Hall et al. [6] studied a wide range of designs spanning diverse platform and mooring system topologies. Therefore, an insight into multiple local extrema (related to different classes of platforms) was essential (an algorithm that can converge to multiple locally optimal configurations simultaneously is very powerful at the initial design stages, where the general characteristics of the design space are studied). The robust Cumulative Multi-Niching Genetic Algorithm (CMNGA) [8] previously developed by the main author was utilised. Its distinctive feature is the ability to preserve the entire population during the entire optimisation process, adding new individuals to the population after each single objective evaluation, never discarding any information. The same algorithm was applied by Benifla and Adam [126] in the FOWT support structure optimisation study. The authors concluded that the CMNGA algorithm took fewer fitness function evaluations and therefore needed less computation time than the standard Genetic Algorithm. An example of design space visualisation of the CMNGA is shown in Figure 2.6.

When working with FOWTs, even if a single class of platforms is considered (e.g., a spar), the design space is likely to be multimodal (containing multiple local optima).

³The objective and constraint violations are combined into one metric

However, since global optimisation generally leads to a significantly increased computational effort, most studies conduct local search instead, leading to less expensive, but still useful, solutions. For instance, Myhr et al. [47] employed the Bound Optimisation by Quadratic Approximation (BOBYQA) algorithm to search for locally optimal taut mooring configurations. This trust-region method⁴, partly depending on the approximation of the objective function, was found to be particularly beneficial in nonlinear, nonconvex problems, with noisy objectives [127]. Therefore, such algorithms can be successfully applied to a problem where the objectives depend on stochastic inputs.

Researchers at the University of Stuttgart [128] and Lemmer et al. [49] used one of the Pattern Search methods. Like for all methods within the direct search optimisation family, the search is only performed in the close vicinity of a starting point, therefore, the results are strongly dependent on the choice of the initial set of variables [80]. For this reason, the local approaches are often complemented by the so-called multi-start strategy, where the optimisation is repeated considering different initial design points (e.g., [51]).

If precision is not of utmost importance, the brute force (or “exhaustive search”) method may be useful. Although it is not the most efficient search method, implemented in the correct way may aid concept-level studies, similar to parametric studies. Brute force optimisation of FOWT support structure was first presented by Sclavounos et al. [45], whose study of a wide range of cylindrical platforms paved the way for the subsequent increasingly complex studies. This optimisation approach has also been applied in [53], following a careful selection of just two free variables, and defining their upper and lower bounds to ensure that the design space was Cartesian (i.e., the ranges of variables do not depend on the values of the other variables).

Derivatives computation

The last consideration is the search method, which refers to the way the optimiser moves from one to the next improved design point. Over the years 2008-2024, studies

⁴Trust-region methods iteratively optimise by approximating the objective within a localised “trust” region and solving a subproblem there, unlike line search methods that move along a direction and select a step length to reduce the objective.

on MDAO of FOWTs have shown a shift towards gradient-based approaches, which have become dominant in the literature since 2020. Prior to this period, stochastic methods were more commonly used. This was mainly driven by the fact that the sensitivity information is not always readily available, and the use of numerical approximation methods (e.g., finite difference) may significantly increase the computation time. The stochastic algorithms applied in FOWT optimisation projects included the previously mentioned BOBYQA, CMN GA and Pattern Search, as well as different versions of the Genetic Algorithm (GA) (e.g., [48]). A comparison of a few derivative-free methods (Pattern Search, Particle Swarm, GA, and Simulated Annealing) in terms of their accuracy and efficiency was performed in [128], showing that the Pattern Search generally requires fewer function evaluations than the other gradient-free methods considered.

Recent (2024-2025) developments in the differentiable boundary element method [129, 130] have significant potential to enable efficient computation of system derivatives in optimisation frameworks that adopt the potential flow solution.

In fact, the gradient-based methods are now preferred due to their higher efficiency, enabling optimisation with respect to a larger number of design variables. For smooth, unimodal and well-behaved problems, these methods outperform the derivative-free techniques in terms of convergence rates [80]. In fact, gradient-based studies tend to incorporate more subsystems into the problem formulation, as discussed in Section 2.2.1.

Within the gradient-based approaches, almost all studies have relied on analytical derivatives facilitated by the differentiable analysis methods employed. These methods, particularly adjoint-based methods, are highly efficient for computing total derivatives. Adjoint methods offer accuracy superior to numerical approximations, such as finite differences or the complex step method, and provide significant improvements in optimisation convergence rates [131]. The computational cost of these methods scales linearly with the number of design objectives and constraints, but remains independent of the number of design variables [132], making them more scalable than the alternative analytical approaches such as direct method [50]. This efficiency enables faster convergence and allows for more complex optimisation problems to be tackled. Up to 2020, optimisation of FOWTs was typically only carried out with respect to a

few design variables. More recent studies utilising the adjoint method show more flexibility, for example, Hegseth et al. [51] optimised a spar-type FOWT using 80 design variables, and Rohrer et al. [56] optimised 36 variables of the main column and tower of a TLP turbine.

In FOWT optimisation studies, the methods using Sequential Quadratic Programming (SQP) are the most common, as this family of methods can handle any degree of nonlinearity. Dou et al. [50] used the gradient-based SQP algorithm implemented in MATLAB Optimisation Toolbox, together with the analytical sensitivities of the mass, damping and stiffness matrices and the equation of motion. Hegseth et al. [51, 52] used the SNOPT (Sparse Nonlinear OPTimiser) algorithm with analytical partial derivatives and total derivatives obtained by the adjoint method [131].

In contrast to the efficient analytical methods, numerical approaches such as finite difference involve calculating partial derivatives of the objective or constraint functions with respect to each design variable and propagating them through the system using the chain rule, which accounts for how each design variable affects the objective (or constraint), both directly and indirectly through intermediate variables. While direct methods can be computationally expensive, they are simpler to implement and do not require the model to be differentiable. However, they scale poorly with the number of design variables, making them less efficient for large-scale problems. These methods are, therefore, not frequently applied in MDAO studies. One exception is the study of Fylling et al. [46], who performed optimisation of 6 variables of the platform, mooring system, and power cable using two-sided finite difference approximations. Although the presented strategy was feasible for improving the design, the authors showed that low accuracy of the numerical partial derivatives might lead to poor convergence behaviour of an optimiser.

It is noteworthy that the gradient accuracy is critical for the success of these methods. Therefore, if the derivatives cannot be evaluated analytically, one of the stochastic optimisation algorithms is usually preferred.

2.3 FOWT numerical analysis approaches

Optimising a complex FOWT system requires analysis methods that are both accurate and efficient, to capture the relevant physics without relying on overly conservative assumptions, thereby enabling more economical designs (Korch, 2024). This section reviews the literature on numerical analysis approaches for FOWTs in MDAO studies, discussing their strengths, limitations and application examples.

2.3.1 Overview of modelling approaches

The literature on FOWT modelling reveals a broad spectrum of tools and techniques, spanning a range of fidelity levels, from high-fidelity computationally demanding time-domain (TD) models to low-fidelity efficient frequency-domain (FD) models. The choice of model depends on the trade-off between accuracy and computational efficiency, as detailed in the following.

Time domain (TD) models TD models include industry-standard software such as OrcaFlex, Sima, and Bladed, and open-source codes such as OpenFAST [133], which can perform fully coupled dynamic nonlinear analyses of the entire floating system. They account for complex interactions between aerodynamic, hydrodynamic, and structural loads and control system under various operating and extreme conditions, including the prediction of the transient responses and nonlinear load-displacement relationships [134].

Despite their extensive capabilities and high accuracy, there are multiple challenges in applying them during the conceptual design. Proprietary TD tools often lack I/O interfaces for integration into broader analysis and optimisation frameworks. Due to the high computational effort, TD models are also impractical in the exploration of a broad design space in the early design phases, as they typically require a wall-clock time of the same order of magnitude (or one order of magnitude less) as the simulated time: for a 10^3 seconds simulation, around 10^2 – 10^3 seconds wall-clock time is required by an average desktop station. While this is not a problem when analysing a single design, it becomes a significant obstacle when analysing thousands of candidate designs

across multiple load cases and random realisations for optimisation purposes.

Even state-of-the-art TD analysis models still face challenges in accurately predicting the response of floating platforms under certain environmental conditions. For instance, mid-fidelity solvers have been shown to underpredict the nonlinear responses of semisubmersible platforms in the low-frequency range [135], even if 2nd-order kinematics and hydrodynamic loads are modelled. The development and validation of high-fidelity tools, such as Computational Fluid Dynamics (CFD) routines, and their application to FOWT analysis problems are currently a focus of intensive research (e.g., [136]).

TD models are typically used during the preliminary and detailed design phases to address specific design questions, building on the insights gained in earlier stages. However, it has been argued that to make more accurate design decisions at the conceptual stage, and thereby reduce costs in later stages, it is crucial to increase the fidelity of analysis from the outset [22]. In line with this, several FOWT MDAO studies applied TD methods in design optimisation, adopting various strategies to reduce computational costs. Advancements in computational capabilities, as well as the use of parallel and cloud computing, have enabled the application of full TD methods in the FOWT conceptual design stage – an example will be given later in this section.

Frequency-domain models FD approaches offer a more efficient alternative for conceptual design, where rapid evaluations are necessary. The flexible, robust, linearised FD models dominate optimisation studies due to their ability to balance accuracy with computational efficiency. They can be used to analyse the linear steady-state response of FOWTs to periodic excitations (e.g., wave and wind forces).

While the FD approaches provide valuable information about the initial design performance, they cannot always be justifiably applied. In operational conditions, especially near rated wind speeds or in extreme environments, nonlinearity of the aerodynamic, hydrodynamic, and mooring loads is non-negligible [137–139], requiring a nonlinear TD solution [139]. Moreover, these methods cannot model transient dynamic effects, such as those experienced in accidental scenarios.

To mitigate these limitations, research has focused on developing extended FD models (or “spectral-domain models” [140]) that aim to derive linear approximations of selected nonlinear effects. Such approaches typically retain the computational efficiency of frequency-domain formulations while incorporating nonlinear behaviour through methods such as equivalent, stochastic or statistical linearisation techniques, or iterative solution procedures in which linearised coefficients are updated based on the response spectrum. Unlike classical linear FD formulations, these methods may therefore require an iterative solution for a given environmental spectrum. The accuracy of these models depends on the accuracy of the linearisation of the true nonlinear behaviour [138].

FD models are often employed in the early stages of FOWT design to explore a wide design space quickly and without the computational burden of TD simulations. They provide valuable insights into system performance but may require more detailed analysis later in the design process.

Several comprehensive multidisciplinary optimisation frameworks based on frequency-domain models have been developed. For instance, Pollini et al. [55] and Ferri et al. [141] employed coupled aero-hydro-elastic dynamic models by incorporating the relevant effects into the linear frequency-domain equations of motion. In the framework developed by Pollini et al. [55], all nonlinear terms were linearised with respect to the undisturbed floater position, which limited the ability to capture the coupling between the loads and the platform’s mean position. Ferri et al. [141] addressed this limitation by linearising the loads around the platform’s steady-state operating point (i.e. the equilibrium position under the steady wind condition). This was achieved by using the linearisation functionality of the time-domain code FAST [133] at each step of the optimisation process. To reduce computational effort, Ferri et al. employed lookup tables for hydrodynamic coefficients. However, the lookup process was based solely on the general dimensions of the platform, assuming an undisturbed underwater shape. As a result, changes in the platform’s underwater shape at the static equilibrium position could not be accounted for.

2.3.2 Time-domain models

As indicated in Tables 2.1 and 2.2, TD models were only applied in a few FOWT MDAO studies. Myhr et al. [47] aimed to minimise the wave load on a Tension Leg Platform (TLP) by introducing a space-frame section in the wave-affected zone. Achieving the required level of accuracy and structural modelling detail necessitated the use of a FEM method together with a time-domain dynamic analysis capable of representing the coupled nonlinear system response. Lemmer et al. [49] performed an optimisation using a reduced-order TD model to study the dynamic behaviour of the turbine coupled with a control system. In a subsequent study [53], the authors compared experimental and numerical TD and FD approaches, concluding that the TLP's response could be sufficiently predicted with the linearised model for preliminary sizing purposes.

More recently, Abbas et al. [59] (2024) performed a series of optimisations of the FOWT structure and controller using the control co-design framework WEIS [142], developed at the National Renewable Energy Laboratory (NREL). The state-of-the-art, open-source, mid-fidelity TD model OpenFAST [133] was used as the primary dynamic model that integrates aerodynamics, hydrodynamics, servodynamics, mooring line dynamics, and structural dynamics into a single, numerically coupled framework. OpenFAST was supplemented with FEM and boundary element method (BEM) models for solving the structural and potential hydrodynamic loads. The FEM model was used to obtain the modal shapes, which were used together with the distributed mass and stiffness properties and nonlinear loads to compute the stress in the tower structures. The Reference Open-Source Controller (ROSCO) [143] was optimised together with the FOWT support structure. The solution of the controller and tower co-optimisation problem, with 19 variables, took nearly two days on a high-performance computing (HPC) machine with 97 CPUs—resources that are typically unavailable to most FOWT developers. Yet, the study co-optimised the controller with either the floating platform, or the tower, but not both. The simultaneous optimisation of the entire system in the time domain, which would be significantly more computationally demanding, has yet to be attempted.

2.3.3 Frequency-domain models

In this section, the background of the FD approach is given first in the scope necessary to understand the differences in specific implementations of the approach in different FOWT optimisation studies.

The formulation of the hydrodynamic model in the frequency-domain frameworks depends on which physical wave-structure interaction mechanisms are considered significant for a given platform configuration and wave regime. Depending on the characteristic dimensions of the platform relative to the wavelength, as well as the importance of viscous effects, different subsets of these load components may dominate the response and motivate different modelling approaches.

The Froude–Krylov force results from integrating the undisturbed incident wave pressure over the wetted surface of the body. This load component depends primarily on the displaced volume and is present even when the structure does not significantly disturb the surrounding wave field. In contrast, diffraction forces arise from the scattering of incident waves by the structure and therefore become important when the characteristic dimensions of the floating members are not negligible relative to the wavelength. Such effects are particularly relevant for large-volume floaters, including spar-type platforms with relatively large cylindrical columns. Radiation forces originate from waves generated by the motion of the floating body itself. These forces are commonly represented through frequency-dependent added mass and radiation damping coefficients, which influence the inertia and damping characteristics of the system. Viscous drag forces arise due to flow separation and viscous effects around the structure and are typically nonlinear functions of the relative velocity between the fluid and the moving body. These loads are particularly important for slender structural members (with a characteristic dimension small relative to the wave height [144]), where they contribute significantly to hydrodynamic damping.

The complete potential-flow solution captures the combined effects of Froude–Krylov loading, diffraction and radiation. In practice, this is typically achieved using boundary element method (BEM) solvers, which require solving one diffraction problem and n radiation problems for each wave frequency and direction. Despite the relatively

high computational cost, this approach has been adopted in many FOWT optimisation studies (for example, in [53] [6]) due to its high physical fidelity. However, the difficulty in obtaining derivatives of the hydrodynamic coefficients has historically limited the application of BEM methods primarily to gradient-free stochastic optimisation approaches.⁵

For some platform configurations, simplifications to the full radiation–diffraction solution can be introduced. For example, platforms with relatively small waterplane areas may exhibit added mass coefficients that remain approximately constant across frequencies, while radiation damping effects remain small compared to inertial forces [146]. In such cases, the added mass may be approximated using analytical two-dimensional coefficients and the radiation damping may be neglected. This simplified representation was employed, for example, by Hegseth et al. [51].

While full diffraction-based solutions provide the highest fidelity among linear hydrodynamic models, intermediate approaches may be adopted when diffraction effects are important but a complete potential-flow solution is computationally undesirable. One such approach is the MacCamy–Fuchs correction, which modifies the inertial term in Morison’s equation using analytical diffraction theory for vertical cylinders. This correction enables partial representation of wave scattering effects while retaining the relative simplicity of Morison-type formulations. It was applied in spar-type FOWT optimisation studies, for example in [51].

If the characteristic dimensions of the structure remain sufficiently small relative to the wavelength, diffraction effects (far-field diffraction/scattering) become negligible, and the hydrodynamic loading may instead be approximated using Morison’s equation. In this approach, the total hydrodynamic load is represented as the sum of inertia and viscous drag contributions. The inertia term accounts for fluid particle acceleration together with hydrodynamic added mass effects associated with body acceleration, while the viscous drag term depends on the relative velocity between the fluid and the moving structure. The relative velocity generally includes contributions from the water particle

⁵A 2024 study by Rohrer et al. [145] demonstrated a first attempt at differentiating the BEM code, which may influence this trend in future optimisation frameworks.

velocity u , current velocity ν_c and platform motion velocity $\dot{\xi}$. Full Morison-based formulations have commonly been applied in FOWT optimisation studies, including, for example, the works of Dou et al. [50] and Myhr et al. [47], where limits on the platform diameter were imposed to ensure the validity of the slender-body assumption. The resulting global equation of motion, obtained by summing the contributions evaluated on individual strips along the structural members, can therefore be expressed as [146]:

$$M\ddot{\xi} + B\dot{\xi} + C\xi = c_m\rho\nabla\dot{u} - (c_m - 1)\rho\nabla\ddot{\xi} + c_d \cdot \frac{1}{2}\rho A_n |\nu_c + u - \dot{\xi}|(\nu_c + u - \dot{\xi}) \quad (2.9)$$

where A_n is the projected area of a submerged structural member per unit length, ρ is the water density, and c_m and c_d are the inertia and drag coefficients, respectively, which depend on the geometry of the floating member.

If viscous effects can be neglected, the drag term on the right-hand side of Equation 2.9 may be omitted, resulting in a linear equation of motion that can be solved directly in the frequency domain. Alternatively, viscous drag effects may be retained within linear frequency-domain formulations through drag linearisation techniques, as discussed later in this section.

As shown in Tables 2.1 and 2.2, first-order potential-flow BEM methods form the foundation of most hydrodynamic modelling frameworks used in FOWT optimisation studies. A gradual progression in modelling fidelity can nevertheless be observed over time. While the hydrodynamic model in [73], one of the earliest optimisation studies, relied solely on linear potential-flow theory, more recent frameworks increasingly supplement the first-order radiation–diffraction solution with additional models accounting for viscous drag effects and, in some cases, second-order wave loading.

For spar-type FOWTs, higher-order wave loads were found to be significantly smaller than the first-order loads, except in the low-frequency (long wave) range [147, 148], or even negligible in terms of the rigid body motions [147, 149, 150]. However, for larger-volume platforms such as semisubmersibles, second-order wave loads may become significant and contribute substantially to low-frequency platform motions and mooring responses. Recent (2023) developments in boundary-element solvers have enabled com-

putation of not only first- but also second-order hydrodynamic loads [151].⁶

2.3.4 Simplified time-domain models

One common simplification used to significantly reduce solution time is downselection of the system's DOFs. Most of the FOWT studies exclude yaw motion, focusing instead on surge, heave, and pitch. This simplification can leverage the symmetry of the platform under analysis but limits the scope to in-plane loads and motion responses. Additionally, most studies neglect the flexible modes, assuming them to be non-critical at the conceptual design stage. However, eigenfrequency analysis is often still performed to ensure that critical system natural periods do not fall within the energetic part of the wave spectrum or coincide with the 1P and 3P rotor excitation ranges [47].

Furthermore, TD and FD codes can model load distributions along selected elastic members, such as the blades or tower, using methods like FEM. Within the TD framework, significant improvements in solution time can be achieved by calculating structural loads and deflections at discrete nodes (e.g., blade root, tower base) instead. This simplification is often implemented using a coupled multibody approach (e.g., [49]), which reduces the number of DOFs by omitting higher-frequency modes.

Another simplification, outlined in [152], involves solving the cross-sectional (internal) loads through a global response analysis in the time domain and identification of the most critical time instances for computing detailed structural responses (e.g., using time-domain FEM). The support structure is simplified and modelled as beam elements. This approach is considered sufficient for early design sizing but has not been demonstrated in design scenarios in the literature up to date.

An alternative solution presented in [152] involves performing a full-time signal structural analysis, with the analysis sped up by either scaling up computational power (using parallel or cloud computing) or reducing the resolution of the structural discretisation (e.g., using a coarser mesh). While the high cost of this solution might be justified in the front-end engineering design (FEED) stage for certification purposes, it

⁶Note that, in general, difference and sum-frequency second-order loads cannot be incorporated straightforwardly into the first-order frequency-domain formulation, in which the structure responds at the same frequencies as the incident waves.

may not be feasible in the conceptual design where a large number of design variants are analysed within a very short time and the structural details are not yet considered.

Significant reductions in simulation time can also be achieved by avoiding unnecessary iterations, recursions, integrations, time domain to frequency domain conversions, and excessive memory usage [153]. Some of these efficiencies can be realised through precomputation of aerodynamic coefficients and mooring line forces. For example, instead of applying the iterative blade element momentum method for aerodynamic analysis during the optimisation, look-up tables of torque and thrust coefficients can be precomputed by running a TD simulation for various tip-speed ratios and blade pitch angles in advance [154]. Similarly, for the mooring system, quasi-static fairlead forces can be derived by interpolating look-up data based on horizontal and vertical displacements as inputs [42]. These simplifications were incorporated into the SLOW code used in the FOWT optimisation study by Lemmer et al. [53], resulting in a high time-to-wall-clock time ratio of 160 (this ratio is typically at least an order of magnitude lower for a full nonlinear TD simulation).

2.3.5 Extended frequency-domain models

FOWT support structure multidisciplinary optimisation frameworks most commonly rely on extended FD analysis models. In these cases, the primary challenge lies in identifying and evaluating accurate linearised approximations of the nonlinear effects that are not captured in the simplest FD models. This section focuses on the improvements to the hydrodynamic, mooring, aerodynamic, elastic, and control submodels.

Hydrodynamic loads

The simplest FD hydrodynamic models typically consider only first-order potential loads, which are much larger than higher-order loads. However, in realistic sea states, load interactions at the difference between two frequencies can excite slow-drift motions, potentially resonating with low natural frequency modes of the floating wind turbine (FOWT) system, such as surge, and increasing dynamic tension in the mooring lines [155]. Similarly, load interactions at the sum of the frequencies can excite the

first tower bending mode [147], significantly influencing fatigue and ultimate loads on taut-moored FOWTs [156].

The analysis of the platform’s response to these loads in the frequency domain is challenging due to the need for linearisation, as well as the fact that the combined first- and second-order wave loads are generally non-Gaussian. Consequently, such analyses are rarely conducted in FOWT optimisation studies. An exception is the work by Lemmer et al. [53], where second-order slow-drift forces were found to be significant for the Triple Spar platform, particularly for low-draft designs. In this study, first-order potential flow solutions were used to derive second-order wave excitation quadratic transfer functions (QTFs). To speed up the calculations, Standing’s Newman approximation was applied, where off-diagonal terms of the QTFs were estimated from the diagonal terms [60,157]. This approximation was considered sufficiently accurate for conceptual design and computationally more efficient than the full QTF approach [158].

The second possible extension, which allows for the accurate modelling of the response of a much wider range of FOWT platforms across a broader range of environmental conditions, is the inclusion of a linearised viscous drag model. While for many platform geometries and load cases the response may be dominated by inertial loads (e.g., large-volume floaters in short waves), the exploration of large design spaces requires a universal approach that is valid for all distinct designs, including those where flow separation effects cannot be neglected. From 2012, FOWT optimisation studies typically incorporate viscous drag in the dynamic models, using two main approaches: the drag term from Morison’s equation (e.g., [6]) and a simpler empirical model (e.g., [128]).

As shown in Equation 2.9, the drag term in Morison’s equation is quadratic with respect to the relative velocity. Several methods for linearising this term in FD approaches have been proposed in the literature, with the stochastic linearisation method by Borgman [159] being the most commonly applied in FOWT multidisciplinary optimisation studies (e.g., [6]). Assuming the wave surface elevation is a zero-mean Gaussian random process and that linear wave theory applies, the fluid velocity is also Gaussian.

This leads to the following expression for the drag load on a thin slice of the platform:

$$\frac{df_d}{dL} = \frac{1}{2}c_d\rho D\sqrt{\frac{8}{\pi}}\sigma_u u \quad (2.10)$$

where σ_u is the standard deviation of the relative velocity, which depends on both the platform's response and the water particle velocity. Since the platform response is needed to calculate the viscous drag, and the drag is an input for response calculations, an iterative approach is required.

Morison's drag term can be split into two components: one contributing to damping (due to platform motion) and the other to excitation (due to wave and current motion). Often, the excitation component is neglected, assuming the water particle velocity is small compared to the platform's velocity under linear wave kinematics (e.g., [48]). However, in some studies, both components are retained and separately accounted for in the equation of motion (e.g., [47,53]). The applicability of the simplified approach can be judged based on the amplitude of motion relative to the size of the structure [160].

Mooring loads

FOWT multidisciplinary optimisation studies based on FD approaches employ various methods to model mooring loads. Most commonly, mooring loads are computed in the time domain and then linearised at each operational point, as seen in studies such as [46,48,50]. This process involves identifying the steady-state (or operational) point, imposing small displacements of the floating structure around this point, and calculating the numerical derivatives of the mooring line loads with respect to displacement (the restoring matrix) using finite differencing. Since mooring loads can be highly nonlinear for large displacements [73], linearisation is typically performed at each operational point, particularly when accurate computation of the mooring line loads is critical. However, in studies limited to platform response modelling and design (e.g., [90]), a simplified approach of linearising around the initial equilibrium position has often been considered acceptable (without proof). The computation of the mooring stiffness (derivatives) can be done either within the optimisation process or in the

preprocessing stage and stored as a look-up table. An example of the latter can be seen in [49].

Alternatively, analytical formulations can be applied, such as the method by Al-Solihat and Nahon [62], which was used in the FOWT optimisation study by Dou et al. [50]. In this approach, the mooring stiffness matrix is derived as a function of the horizontal and vertical projections of the mooring line, the vertical position and radius of the fairleads, and the horizontal and vertical forces at the fairlead, which are obtained using a catenary equation.

Another unique solution was proposed by Larsen and Sandvik [161] and later applied in an optimisation framework by Hegseth et al. [51]. In this method, the mooring lines are modelled as a single-degree-of-freedom spring-mass-damper system. The inertia is derived from the quasi-static configuration of the mooring line, while the stiffness is computed based on the displacement of the line from equilibrium to a point of significant motion (twice the standard deviation of platform motion). The damping is derived from the quasi-static configuration of the mooring line using stochastic linearisation of a variant of Morison's equation. Unlike most FOWT optimisation studies, this approach includes both the inertia and damping of the mooring lines in the FD equation of motion, alongside the mooring stiffness.

Finally, a frequency-domain lumped-mass model has recently (2025) been proposed by Abdelmoteleb and Bachynski-Polić [162]. Unlike the previously mentioned approaches, this model was able to estimate the wave-frequency dynamic tension component, with good agreement with fully nonlinear coupled time-domain simulations across multiple mooring designs and loading conditions.

Aerodynamic loads

Although the rotor and support structure designs are typically optimised separately [163], it is crucial to consider aerodynamic loads when optimising the support structure. The aero-elastic loads from the rotor generate complex, nonlinear, time-dependent forces and moments on the foundation [164]. While aerodynamic loading is primarily determined by the rotor design, it is also influenced by the motion of the rotor, which in

turn depends on the design of the support structure.

Considering the standard frequency-domain equation of motion:

$$[-\omega^2 (\mathbf{M} + \mathbf{A}(\omega)) + i\omega\mathbf{B}(\omega) + \mathbf{C}] \boldsymbol{\xi}(\omega) = \mathbf{F}(\omega) \quad (2.11)$$

where $\mathbf{M}$, $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are the 6×6 mass, added mass, damping, and hydrostatic stiffness matrices, $\boldsymbol{\xi}$ is the complex response amplitude vector, and $\mathbf{F}$ the excitation force, the fluctuating aerodynamic loads can be accounted for by incorporating the contributions of wind excitation (on the fixed structure) on the right-hand side, and the aerodynamic added mass, damping, and stiffness (related to the motion of the structure) can be contribute to the terms on the left-hand side. The added mass represents the additional inertia imparted to the system by the acceleration of air around the moving rotor-nacelle assembly. Aerodynamic stiffness arises because when the platform pitches, the orientation of the rotor plane changes relative to the average wind direction, altering the induction factors and, consequently, the aerodynamic loads. Aerodynamic damping, on the other hand, accounts for changes in aerodynamic loads due to variations in the relative wind speed as the platform moves.

Most of the published literature on FOWT optimisation relies on steady-state solutions from the standard BEM approach. Two-dimensional lookup tables for power and thrust coefficients are created at different tip-speed ratios and blade pitch angles. Since aerodynamic torque and thrust also depend on the square of the rotor's effective relative wind speed, linearisation is required before integrating these results into a FD framework. This is typically done by expanding the standard thrust and torque formulations using a first-order Taylor series (tangent linearisation) [165]. Abdelmoteleb and Bachynski-Polić provided a comprehensive comparison of the frequency-domain aerodynamic models in [166]. In addition, various semi-empirical corrections, such as those for unsteady inflow, can be applied, as demonstrated in [47].

Hegseth et al. [51] employed a slightly different approach by applying linearised blade element momentum theory with semi-empirical corrections. Instead of precomputing thrust and torque coefficients, the authors directly calculated the normal and

tangential induction factors for each blade element. The resulting aerodynamic forces on the rotor (thrust and torque), as well as the blade root loads, were then obtained by integrating the normal and tangential forces along the blade’s length. These loads were subsequently linearised using a Taylor series expansion. To compute the partial derivatives of the loads with respect to wind speed, blade pitch angle, and rotor speed (gain factors [167])⁷, numerical perturbations were applied to these inputs, and the resulting changes in forces were observed. Since the linearisation was performed for each blade element, detailed information about the load distribution along the blade was readily accessible.

The method used in [6, 48, 73] involves computing aerodynamic matrices using the linearisation capabilities of TD software. In these studies, the added mass, stiffness, and damping coefficients were determined at each wind speed condition at zero pitch angle. The linearisation assumes that the platform’s pitch motion about the equilibrium position is relatively small, and that the blade pitch angle is held fixed at the value required for the given mean undisturbed wind speed. These loads can either be computed directly during the optimisation process or precomputed, as was done, for example, in [55].

Structural dynamics

Up until 2020, most FD models in the reviewed FOWT optimisation studies employed the rigid body assumption, with structural dynamics rarely included. One early exception was the rigid multibody system approach used in [49], where an additional DOF was introduced to account for tower-top fore-aft displacement due to deformation. This addition allowed for a simplified representation of tower flexibility. A similar method was applied by Hegseth et al. [51] (referred to as the “generalised displacements approach”), who noted that this method worked well for the tower but was too simplistic to accurately capture the blade deflections.

Recent (post 2021) studies present a mix of approaches, with the rigid body assump-

⁷In control systems, gain factors are numerical values that represent how much an output changes in response to a change in an input.

tion still present in some (e.g. [55], [58]), but with a growing dominance of multi-body or generalised displacement approaches. Notably, two recent studies [57] and [59] introduced high-fidelity FEM models—an approach that had not been seen in FOWT MDAO studies before 2023, except for the early work of Myhr et al. [47]. In [57], the FEM model was employed for eigenvalue analysis. The TLP FOWT structure was represented using three-dimensional Euler-Bernoulli hollow beam elements with defined mass and stiffness properties. The pontoons, central column, and rotor-nacelle assembly were connected through rigid links, and point masses were added to represent ballast and the rotor-nacelle assembly. Linear point stiffnesses were applied to account for hydrostatic and mooring stiffness.

Control system dynamics

Modern utility-scale horizontal-axis wind turbines (HAWTs) typically employ a variable-speed, variable-pitch control strategy. As illustrated in Figure 2.7, in low wind speeds, the rotor speed is adjusted to match changes in wind conditions and maintain the optimum tip-speed ratio (TSR), while in above-rated wind speeds, the rotor blades are pitched to maintain the rated power output and protect the turbine from potential damage. Incorporating this control strategy into an analysis tool requires the addition of two extra DOFs: rotor speed and blade pitch angle. However, modelling the dynamics of a multidisciplinary system, including the control system, can be computationally demanding and requires considerable effort in implementation. Consequently, this aspect received limited attention in the early FOWT optimisation literature until relatively recently (around 2020).

The integration of control system dynamics into FOWT optimisation studies gained prominence with the work of Lemmer et al. [49], who highlighted its significant impact on power production performance. In this study, reduced-order nonlinear TD simulations were initially used to determine the system’s steady state. From this steady-state solution, a linearised model was derived and utilised to design an optimal model-based Proportional-Integral (PI) controller for each support structure design iteration. The study implemented the baseline controller of the NREL 5 MW reference wind tur-




Figure 2.7: IEA 15 MW) controller regulation trajectory – variable speed variable pitch strategy. Adapted from [9]

bine [43] for below-rated wind speeds and a Linear Quadratic Regulator (LQR) with state feedback controller for above-rated wind speeds. The control outputs included blade pitch angle and generator torque, and the inputs included the states of the linearised model—platform surge, heave, and pitch displacements, tower-top fore-aft displacement, rotor speed, and their derivatives. At the end of each optimisation iteration, a nonlinear TD code was used to simulate the system response and evaluate the design objectives. It is important to note that this approach was sequential, with the control system being tuned independently of the other design variables.

In a 2020 optimisation study by Hegseth et al. [51], the control system was modelled as a first-order state-space system, with rotor speed as the input and generator torque and blade pitch angle as the outputs. To prevent high-frequency excitation of the controller, a low-pass filter was applied. For low wind speeds, generator torque was set proportional to the square of the rotor speed, while for higher wind speeds, the blade pitch angle was adjusted using a PI control strategy. The structural and control models were integrated into a single closed-loop system, which was then transformed into the frequency domain for response calculation. This approach allowed for a simplified yet sufficiently accurate coupled analysis of the entire system.

A similar approach was adopted in the later work of Hegseth et al. [52], where four control strategies were compared. These included a basic gain-scheduled PI controller and three variations that incorporated additional platform pitch/nacelle velocity feed-

back and a low-pass filter. The introduction of velocity feedback helped reduce platform pitch motion and fatigue loads in the tower, while the low-pass filter improved the power output quality.

In [53], a Single-Input-Single-Output (SISO) PI controller was implemented, with generator torque and blade pitch as the outputs to the plant, and rotor speed and tower top displacement as inputs, influenced by wind and wave disturbances. The system and disturbance transfer functions were derived from a linear model of the dynamics, similar to the methods in the previously mentioned studies. Additionally, it was observed that the viscous drag coefficient, which depends on the Keulegan-Carpenter (KC) number and, consequently, on the relative flow velocity and platform response, plays a significant role in the tuning of the controller gains. As a result, both the viscous drag and controller gains were computed iteratively to account for these interactions.

2.3.6 Multifidelity models

The traditional design approach has evolved into a multi-fidelity framework, where low-, medium-, and high-fidelity tools are used concurrently to evaluate concepts [18]. One common strategy is to employ models of increasing fidelity at successive stages of the concept design process. For instance, Hall et al. [168] proposed an approach where a simplified linearised model is used for wind farm layout optimisation, a quasi-static analysis for the optimisation of mooring lines, and a TD coupled analysis for deriving correction factors (dynamic/static load ratios) applied iteratively to the less accurate models.

While formal multi-fidelity optimisation approaches have not yet been widely applied to the design of FOWTs, they offer a promising solution. These methods construct approximative or reduced-order models at chosen stages of the optimisation process to improve computational efficiency. For example, Jasa et al. [169] conducted a surrogate-based optimisation (SBO) for bottom-fixed wind turbines, using both low- and high-fidelity models at initial design points to derive correction functions between the models for each function of interest (objectives and constraints).

2.3.7 Machine learning models

Although high- and mid-fidelity physics-based models are the most accurate, they are computationally expensive and thus not ideal (or even feasible) for highly iterative tasks such as design optimisation. Alternatively, data-driven models offer similar accuracy to nonlinear dynamic models but at a significantly reduced computational cost.

In a typical supervised-learning setting, a surrogate is trained offline on a database of input-output pairs generated using higher-fidelity numerical simulations or experiments, thereby learning an approximation of the underlying input-response mapping. Once trained, the surrogate can be evaluated in milliseconds, enabling its use within gradient-based or gradient-free optimisation, uncertainty quantification, and reliability analyses that require thousands to millions of model evaluations [170, 171]. In this way, the computational burden is shifted from the optimisation loop to an upfront pre-computation phase devoted to dataset generation and model training.

The benefits of using surrogate models prior to FOWT optimisation were explored in [128]. It was observed that, to properly interpret the optimisation results, a comprehensive understanding of the overall system behaviour is essential. The surrogate model served this purpose and was subsequently used to test and evaluate the optimisation algorithm before applying it with the actual simulation tools. Surrogate models can also be employed for design space screening (i.e., identifying the most critical design variables) and uncertainty assessment.

Surrogate models can also serve as inexpensive stand-ins for optimisation studies (SBO [172]). This approach shifts the computational effort from the optimisation process itself to the preprocessing stage. Instead of running time-consuming simulations during the optimisation, approximate models are evaluated at each iteration, significantly reducing computational time. The effectiveness of this approach has been demonstrated in two frameworks for the optimisation of TLP [173] and semisubmersible [174] offshore oil and gas (O&G) floating structures. In both cases, the most probable maximum heave motion was modelled as a response to five hull shape variables. Artificial Neural Networks (ANNs) and Radial Basis Function (RBF) models were trained using approximately 700 samples obtained through FD analysis. Additionally, Pillai et

al. [89] applied surrogate-based optimisation to offshore renewable energy mooring systems. The authors used a Genetic Algorithm as the optimiser, querying Random Forest regression and classification models to evaluate the design objectives (cumulative fatigue damage and material cost) and assess the satisfaction of constraints (mooring tension). These models were trained based on OrcaFlex simulation results for various combinations of mooring parameters and different anticipated sea states.

A similar approach was demonstrated by Li and Choung [175], who trained a multilayer feedforward ANN to predict wide-banded fatigue damage in catenary mooring lines of a FOWT. Environmental variables (significant wave height, zero-crossing period, wave direction, and current speed) served as inputs, while tension range distributions were the outputs. The trained model showed good agreement with time-domain simulations, though the authors noted that small inaccuracies in predicted stress ranges can lead to non-negligible errors in fatigue damage estimates due to the nonlinear relationship.

In the context of FOWT optimisation, Zhang et al. [176] conducted a multi-objective optimisation of a semisubmersible platform using a Standard Response Surface, which was constructed from a database of results obtained from OpenFAST. In this study, the platform parameters served as inputs, while the outputs were the motions and natural frequencies, which also acted as design constraints. Trubat et al. [177] performed FOWT mooring optimisation using a Differential Evolution algorithm, with an Artificial Neural Network (ANN) model predicting static and dynamic responses (line forces and platform motions) based on input parameters related to line geometry. Ye et al. [178] optimised the FOWT mooring system using the NSGA-II algorithm and a Kriging model, which was trained on results from OpenFAST. As in the previous studies, the mooring system parameters were input to the model, while platform motions (which also served as design constraints) were the outputs. Wang et al. [179] trained a Deep Neural Network (DNN) using the geometric parameters of the mooring system of a semisubmersible FOWT to predict optimisation objectives – mooring tension and platform motions – for three design load cases. A key innovation in the approach was using initial static calculations to filter out infeasible samples from the dataset, which

were then excluded from the training of the surrogate model.

Unlike the other studies, a semisubmersible FOWT optimisation by Bessone et al. [180] employed a surrogate model for potential hydrodynamic coefficients (rather than directly modelling the design constraints). In this case, the data-driven replacement of the boundary element method code was used within a highly efficient frequency-domain framework to compute system responses, which were then considered design constraints. The authors trained an XGBoost model based on results from the HAMS [181] hydrodynamic solver, implemented with the RAFT code [182]. A similar approach was presented by Rohrer et al. [183], where polynomial surrogate models trained within the OpenMDAO framework were used to predict both the potential coefficients and their exact derivatives, thereby enabling gradient-based optimisation of semisubmersible platforms across multiple geographical locations. Earlier studies by Rohrer et al. [145] and Coraddu et al. [184] also developed surrogate models for hydrodynamic coefficients, although these works did not directly apply the models to design optimisation.

In fact, an earlier study by Kalimeris et al. [185] investigated surrogate-modelling strategies for semisubmersible FOWT optimisation, including both “end-to-end” surrogate replacement of the full analysis workflow and selective replacement of the computationally expensive hydrodynamic solver (PyHams) within the RAFT framework. The authors concluded that the PyHams-surrogate approach achieved a better accuracy, while the “end-to-end” approach was more computationally efficient.

All the studies discussed above employed surrogate models to accelerate design evaluations within stochastic optimisation frameworks. With only two exceptions, these studies used surrogate models to predict design objectives and constraints, rather than replacing specific bottleneck models in the coupled solution process. Furthermore, each study trained its models for specific environmental conditions, mooring studies focused on particular platform designs, and platform studies were based on specific mooring configurations. Despite the growing use of data-driven models in FOWT analysis and design, surrogate models have not yet been applied in multidisciplinary optimisation approaches.

2.4 Summary and research gaps

This chapter provided background information and reviewed the literature and best practices concerning the key building blocks of FOWT optimisation: numerical analysis, design problem formulation, and optimisation algorithms. A particular focus was placed on MDAO as an emerging methodology for optimising complex coupled engineering systems. Moreover, the concept of ICR was introduced, addressing the common misconceptions surrounding it. This laid the foundation for a deeper understanding of the ICR concept and its application in the design process, as will be explored in the final technical chapter of this thesis.

The literature review highlighted a gap in numerical dynamics modelling, where optimisation studies typically rely on either reduced-accuracy, fast dynamic models to cover a broad design space for holistic design optimisation, or higher-accuracy, nonlinear models focused on the design of specific components. There is a need for a model that balances sufficient accuracy across a wide design space with low computational cost and solution time, enabling optimisation with respect to a comprehensive set of design variables across the entire FOWT system, while adhering to realistic design constraints. Multifidelity approaches are not yet commonly applied, although first studies on surrogate-based FOWT optimisation have emerged.⁸

In that respect, current FOWT data-driven surrogates are primarily used to predict design objectives and constraints, not the dynamic loads and responses. For example, they predict mooring line fatigue damage without insight into system motions or mooring tensions. Directly linking design variables to constraints, masking the underlying dynamics, reduces the interpretability of the analysis. Such models face industry reluctance due to the non-transparent solution and lower confidence in results. Moreover, these models have so far been trained on datasets specific to particular loading conditions and floating platform/mooring subsystem designs. This limits their generalisability and hinders integrated system-level optimisation. No data-driven approaches have

⁸The WEIS framework, currently under development at NREL [142], enables multifidelity dynamics modelling within an optimisation context. However, it has only recently been made publicly available (after many of the methods of this thesis have already been selected/developed).

been applied in FOWT MDAO studies to date.

While the global dynamic response of FOWTs to environmental loads is relatively well understood, and numerous models have been developed to represent this behaviour, the identification and role of the ICR have been largely overlooked. This gap persists despite the ICR's potential significance in characterising and interpreting the system's complex motions. Misconceptions regarding the location of the ICR remain widespread, contributing to its continued neglect. A deeper understanding of the ICR under realistic operational conditions could enable the design of FOWTs that leverage their inherent floating characteristics, rather than focusing solely on minimising platform motions.

The main gap identified in the formulation of the design problem lies in setting constraints based on motion response rather than directly addressing the structural ULS and FLS loads. In particular, the fatigue of mooring lines has been largely overlooked as a design constraint in FOWT MDAO studies. While this approach may result in conservative designs, it can also lead to structurally unsafe outcomes, necessitating costly revisions to the conceptual design once structural checks are performed during the detailed design stage. Additionally, natural frequency constraints are always set based on values corresponding to the peak exciting frequencies of environmental factors (such as wind and waves) and rotor frequencies, with an arbitrarily applied margin. To date, no approach has been presented that considers the probability of these exciting frequencies and the energy contained within the relevant frequency bandwidth.

Overall, there is a positive trend away from sequential component-level design and brute-force optimisation towards integrated system-level optimisation, though challenges remain. In this context, the research presented in this thesis contributes to advancing the complexity and realism of FOWT conceptual design optimisation, which is key to accelerating the design and deployment of economically competitive FOWTs.

Chapter 3

Multidisciplinary design analysis optimisation framework

This chapter is based on the published paper:

P2 Patryniak K, Collu M, Coraddu A, Multidisciplinary Design Analysis and Optimisation of Floating Offshore Wind Turbine Support Structures - Coupled Model Solution Strategies (2023). *Journal of Physics Conference Series 2626*.

3.1 Introduction

Building upon the coupled-system challenges and MDAO motivation introduced in Chapters 1 and 2, this chapter focuses on the implementation of the baseline coupled frequency-domain framework and the treatment of cyclic subsystem interactions relevant to FOWT optimisation. Specifically, this chapter aims to investigate:

- the importance of a coupled dynamic model for accurately computing the response of FOWTs and quantifying design constraints, and
- the influence of the choice of iterative solution strategy on the efficiency of the coupled analyses.

3.2 MDAO framework

This section outlines the MDAO framework, starting with the global coupled model and its architecture, which defines the flow of information between the aero-hydrodynamic subdisciplines (or components). It then provides details of the individual disciplinary models and groups of models (cyclic subsystems).

3.2.1 Coupled model

The global coupled model was constructed with disciplinary submodels connected through input-output relationships, as shown in the N-squared diagrams [186] in Figures 3.1 and 3.2, following the extended design structure matrix (XDSM) methodology [5], and leveraging the functionalities of the open-source OpenMDAO package [79].




Figure 3.1: Coupled model Design Structure Matrix - Optimisation loop.

The multidisciplinary feasible (MDF) architecture has been selected to build the framework, where the optimiser (dark grey block) controls the design variables (x), while the coupling variables exchanged between the disciplinary models ($\hat{u}$) are directly managed by the coupled model (MDA) solver (light grey block), as seen in Figure 3.1. For a given design “proposed” by the optimiser, the solver attached to the MDA group finds a set of internal component states and coupling variables consistent across the model, i.e., converges the coupled model by driving the total system residual ($\hat{r}$) to zero. Once the multidisciplinary analysis has converged, the relevant outputs are passed to a single component (red block) that evaluates the design objectives and constraints, which




Figure 3.2: Coupled model - Design Structure Matrix - Multidisciplinary Analysis (MDA) loop.

then inform the optimisation process. For the MDF architecture, the optimisation problem is therefore formulated as follows:

$$\begin{aligned}
 &\text{Minimise } f(x; \hat{u}^*) \\
 &\text{by varying } x \\
 &\text{subject to } g(x, \hat{u}^*) \leq 0 \\
 &\text{while solving } \hat{r}(\hat{u}; x) = 0 \\
 &\text{for } \hat{u}.
 \end{aligned} \tag{3.1}$$

where $\hat{u}^*$ is a multidisciplinary feasible point (set of coupling variables) found by the MDA and available to the optimiser at each iteration.

Since the focus of the study covered in this chapter was on developing a coupled dynamic model, with optimisation strategies not yet considered, this simple architecture was deemed most appropriate. Still, the model was constructed in a way that facilitates integration in an optimisation framework.

Figure 3.2 provides a closer view of the MDA group. The green boxes represent the disciplinary components (analysis tools), while the parallelograms denote the design

variables and coupling variables that are passed between the components. Inputs are passed vertically, while outputs are passed horizontally. Variables above the diagonal represent feedforward interactions, whereas those below the diagonal represent feedback interactions.

For example, the damping module receives information from the potential problem module (potential damping B_h), the aerodynamics module (aerodynamic damping B_{aero}), and the viscous flow module (viscous damping B_v), and combines them to compute the total damping B . This total damping is then passed forward to the RAO module. The information “flows” through the statistics and viscous flow analysis blocks (components), ultimately returned to the damping module via the viscous damping coefficient B_v . Such feedback connections enable the modelling of two-way coupling between the selected components, forming two groups, or cycles, as illustrated in two shaded regions of Figure 3.2.

3.2.2 Fundamental disciplinary model

The FOWT support structure is assumed to be rigid, as justified by the insignificant effect of elasticity on the wave-frequency motion response of a spar platform [187].¹ Since the dominant aerodynamic and hydrodynamic loads act primarily in the longitudinal direction, the response of the spar-type FOWT is dominated by the in-plane DOFs: surge, heave and pitch, while out-of-plane motions remain comparatively small. Hence, only in-plane DOFs are considered here. The model is based on solving the static equilibrium under the influence of mean wind and mooring loads, followed by a dynamic analysis at this equilibrium position. The models utilised to solve the static and dynamic responses are summarised in Table 3.1 and detailed in the following sections.

¹The elasticity of the platform was found to have a significant impact on the 1st bending mode frequency range tower base fore-aft bending moment, but not on platform motions or wave-frequency loads.

Table 3.1: Summary of modelling approaches within the analysis framework with reference to XDSM blocks.

Component	Name in XDSM	Modelling approach
Mean aerodynamic load; Aerodynamic damping	<i>Aero</i>	Precomputed BEM with corrections; single wind direction
Mean mooring load; Mooring stiffness	<i>Mooring</i>	Quasi-static model (MoorPy)
Inertia matrix	<i>Inertia</i>	Analytical
Hydrostatic stiffness matrix	<i>Hydrostatics</i>	Hydrostatic pressure integration (Capytaine)
Gravitational stiffness matrix	<i>Restoring</i>	Analytical
1st order potential hydrodynamic load	<i>Potential</i>	Boundary-element method (Capytaine); single wave direction
Viscous hydrodynamic load	<i>Viscous</i>	Borgman's linearisation of the Morison's equation
Control	N/A	Fixed point from control table

Static equilibrium

The static equilibrium of the system is defined through Equation 6.1:

$$(\mathbf{C}_h + \mathbf{C}_g)\bar{\boldsymbol{\eta}} = \bar{\mathbf{F}}_{aero}(\bar{\boldsymbol{\eta}}) + \bar{\mathbf{F}}_{moor}(\bar{\boldsymbol{\eta}}) \quad (3.2)$$

$\bar{\boldsymbol{\eta}}$ is the mean offset of the platform, $\mathbf{C}_h$ and $\mathbf{C}_g$ are the hydrostatic and gravitational stiffness, and $\bar{\mathbf{F}}_{aero}$ and $\bar{\mathbf{F}}_{moor}$ are the mean aerodynamic and mooring loads on the FOWT.

The mean aerodynamic load was pre-computed for a range of mean wind speeds and tower inclination angles using the aerodynamic time-domain code AeroDyn [188], which employs blade-element momentum theory (BEMT) along with the Pitt and Peters skewed wake correction model, as well as Prandtl tip-loss and hub-loss corrections. The simulations were run in steady wind conditions with a power law exponent of 0.2.²

²This used to be the typically assumed exponent, recommended by the IEC 61400-1 standard [189] for the normal wind model; however, later studies showed that it is only accurate for a very small set of

Blade pitch angle and generator speed were fixed for each simulation at an appropriate operational point according to the control table, dependent on the mean hub height wind speed. The wind speed was sampled at 1 m s^{-1} while the inclination angle was sampled at 0.5° . The aerodynamic mean load was obtained by linear interpolation from the look-up table. The precomputed aerodynamic mean load is presented in Figure A.1 in Appendix A alongside the control parameters.

The total load on the platform from the mooring system was computed using the quasi-static mooring model MoorPy [190]. This model determines the instantaneous position and tension in the mooring lines based on the net weight and axial stiffness, employing standard catenary equations. The analytical formulation of these nonlinear catenary equations enables efficient computation of the mooring stiffness matrix. MoorPy is used in two stages of the overall solution: i) to obtain the total load on the platform from the three mooring lines as part of the static equilibrium solution, and ii) to calculate the tension in the mooring lines and derive the mooring stiffness matrix at the equilibrium position as part of the dynamic solution.

Since both the aerodynamic and mooring loads need to be computed at the mean offset of the platform, and the calculation of the mean offset depends on these loads, an iterative solution approach must be adopted. The *Equilibrium*, *Aero*, and *Mooring* components in Figure 3.2 form a *group* (cycle 1 in Figure 3.2) (*Equilibrium* component is introduced to define the static equilibrium of Equation 6.1). The states of the three components: $\bar{\eta}$, $\bar{\mathbf{F}}_{aero}$, and $\bar{\mathbf{F}}_{moor}$ are determined iteratively, following one of the strategies outlined in Section 3.3.

This approach enables the updating of load values based on both the rotor plane inclination and the platform offset. Thus, the two-way coupling between the aerodynamic and mooring loads and the attitude and position of the floating system is effectively captured.

atmospheric and marine conditions, and currently lower values 0.11 - 0.14 are more commonly applied for offshore conditions. The exact value does not impact the discussion made in this thesis.

Dynamic response

The dynamic response of the platform at the static equilibrium position is solved using a linear frequency-domain approach, where the platform's response to wave excitation is computed based on the solution to the forced oscillation problem. The complex amplitude of motion in the j - th direction, resulting from excitation in the i - th direction, is given by the following Equation 3.3 [191]:

$$\sum_{j=1}^n \eta_j [-\omega^2 (M_{ij} + \mathbf{A}_{ij}) + i\omega \mathbf{B}_{ij} + C_{ij}] = a \mathbf{X}_i \quad (3.3)$$


where a is the wave amplitude, M_{ij} is the system mass (inertia), C_{ij} is the stiffness (sum of mooring, gravitational, and hydrostatic stiffness), $\mathbf{A}_{ij}$ is the frequency-dependent added mass, $\mathbf{B}_{ij}$ is the damping (sum of the frequency-dependent potential hydrodynamic radiation damping, viscous hydrodynamic damping, aerodynamic damping, and an additional linear damping, and $\mathbf{X}_i$ is the frequency- and direction-dependent first order complex wave excitation per unit wave amplitude. The additional linear damping is included as an artificial correction term to reconcile the numerical model with experimentally measured responses of the original OC3 design [149], which underestimated damping relative to tank tests. The same correction is retained across all platforms analysed in this thesis.³

The inertia and gravitational restoring matrices are computed using standard analytical formulations (see [192], [191]). Hydrostatic stiffness is calculated by an external module Capytaine [193] by integrating the hydrostatic load on the submerged part of the floating platform. Mooring stiffness is determined using the quasi-static mooring model MoorPy (as detailed earlier in this section).

The frequency-dependent potential hydrodynamic coefficients (added mass, radiation damping, and excitation⁴) were computed using a 3D linear frequency-domain

³In the optimisation context, this may inadvertently favour platforms with lower inherent damping (e.g., slender geometries), for which the fixed correction term partially compensates for the damping they would not physically have.

⁴The transfer function between the wave amplitude and the sum of Froude-Krylov and diffraction loads on the platform

Figure 3.3: Panels convergence study. d -normalised panels density.

Boundary element method (BEM) solver, Capytaine [193],⁵ for the floating platform in the equilibrium position.

The underwater surface of the floating body was discretised using the Capytaine mesher, with an axisymmetric surface defined by rotating a parametrically defined profile curve around the vertical axis. Quadrilateral panels were created for the lateral and bottom surfaces. To ensure adequate mesh refinement, the Nyquist-Shannon sampling theorem [194] was followed, enforcing a minimum of four panels per wavelength. A grid convergence study was also performed, computing the surge dynamic response with increasing panel densities. As shown in Figure 3.3, the result converges to the higher-fidelity model as panel density increases, but at the cost of higher computation time. The selected mesh density balanced accuracy and efficiency, with further resolution increase yielding no significant accuracy improvements.

This approach improves upon those presented in the literature by accounting for the modified underwater geometry of the floating platform at its equilibrium position when computing the hydrodynamic loads.

⁵Capytaine is a Python implementation of the code Nemoh developed at the École Centrale de Nantes

Viscous loads are reintroduced by adding a viscous damping coefficient to other damping sources (*Viscous* and *Damping* components in Figure 3.2). This is achieved through a stochastic linearisation of the drag term in Morison's equation, as proposed by Borgman in [159]. For each horizontal slice of the floater (dL) the viscous drag (dF_d) is given by:

$$dF_d/dL = 1/2 \cdot c_D \rho D \sqrt{8/\pi} \sigma_u u , \quad (3.4)$$

dependent on the drag coefficient (c_D), characteristic dimension (D), relative normal velocity u and its standard deviation (σ_u). By definition, the standard deviation of a response velocity is calculated as the square root of the zeroth moment of its response spectrum (in the *Statistics* component in Figure 3.2):

$$\sigma_u = \left(\int_0^\infty S_u(\omega) d\omega \right)^{1/2} \quad (3.5)$$

The velocity response spectrum is obtained as a product of the square of the velocity RAO and the sea spectrum (JONSWAP):

$$S_u(\omega) = |RAO_u(\omega)|^2 S_\zeta(\omega) \quad (3.6)$$

The velocity RAO is obtained by differentiating the displacement response in the frequency domain:

$$RAO_u = \omega \cdot RAO_\eta \quad (3.7)$$

The displacement RAO is defined as the motion response amplitude (translated to each horizontal strip) divided by the wave amplitude (*RAO* component in Figure 3.2):

$$RAO_\eta = \eta/a \quad (3.8)$$

In general, the relative normal velocity contains contributions from both the wave-induced water particle velocity and the platform motion velocity. In the present formulation, linear wave kinematics are assumed and the water particle velocity contribution

is neglected relative to the platform motion in order to simplify the solution.⁶ Viscous damping is then computed according to its definition:

$$B_v = dF_d/du \quad (3.9)$$

Since the drag load given by this approximate expression depends on the platform's response statistics (specifically, the standard deviation of velocity due to platform motion, σ_u in Equation 3.4), and the response calculation requires viscous drag as input, viscous damping and motion response form a feedback cycle (*cycle 2* in Figure 3.2).

The dynamic motion of the FOWT affects both the speed and direction of the relative inflow to the rotor, thereby influencing the aerodynamic loads. To account for this, the linearised aerodynamic damping matrix was introduced and added to the remaining sources of damping in B of Equation 3.3. Aerodynamic damping was precomputed with AeroDyn for a range of wind speeds and tower inclination angles. The wind speed was varied in small increments to numerically compute the derivatives of the aerodynamic load with respect to wind speed, i.e., damping [195]. The resulting coefficients were stored in a lookup table for use in the framework. For more details and results, see Appendix A. The turbulent wind excitation was not considered in this fundamental model. However, it will be considered later in Chapter 6.

The verification study of the present model is reported in the underlying publication [196] and is omitted here, since the verification of the more advanced models developed in Chapter 6 provides a more comprehensive assessment of the modelling framework.

3.3 System solution strategies

The coupled model consists of two cyclic subsystems. Cycle 1 aims to converge the mean loads and the mean position of the FOWT, while Cycle 2 iteratively solves for the viscous damping and the velocity of the floating structure. These subsystems require an iterative solution. Several iterative solvers are considered, differing in how

⁶Similar approach was adapted for example in [6, 48].

the coupling variables (i.e., responses exchanged between components) are handled.

In the following, each solver is introduced briefly, with algorithms and XDSM diagrams available in Appendix A. A comparative overview is given in Table 3.2. Additionally, numerical examples illustrating the performance of each solver are presented in Appendix B. At the end of this section, a demonstration of one of the solvers' implementations in the current MDAO framework is given.

Table 3.2: Comparison of nonlinear solvers in terms of state updates and convergence behaviour.

Solver	States	Convergence
Block Jacobi	Uses previous iteration values	Linear, slow
Block Gauss-Seidel	Uses most recent values as they are computed	Linear, faster than Jacobi
Block Gauss-Seidel with Aitken relaxation	Uses most recent values with Aitken acceleration	Linear, accelerated convergence
Purely monolithic Newton	Solves the full coupled system simultaneously	Quadratic near the solution
Full-space hierarchical Newton	Performs inner subsystem solves before the global Newton update	Quadratic, robust and faster for strongly coupled systems
Broyden	Maintains and updates an approximate Jacobian; uses the latest states	Superlinear

3.3.1 Iterative solvers

The **Nonlinear Block Jacobi solver** (NLBJ) [197] begins with an initial guess for all coupling variables ($\hat{u}^{(0)}$). It repeatedly updates all coupling variables ($\hat{u}_i$) across the components simultaneously, until convergence is achieved, as outlined in Algorithm 1 in Appendix A, reproduced from [11], and illustrated in Figure A.1. Since each component receives input from the previous iteration, the components can be run in parallel without the need for communication between them.

The **Nonlinear Block Gauss-Seidel solver** (NLBGS) [198] solves each component sequentially, using the most up-to-date values of the coupling variables available.

As a result, it typically converges in fewer iterations than the Jacobi solver. However, because the components must be executed in sequence, it cannot take advantage of parallel computing, resulting in each iteration taking longer.⁷

The **Nonlinear Block Gauss-Seidel solver with Aitken relaxation (NL-BGS+A)** [199] updates the coupling variables based on the latest values available modified by a relaxation factor θ , as described in Equation 3.10:

$$\hat{u}_i^{(k)} = \hat{u}_{temp} + \theta \Delta \hat{u}_i^{(k)} \quad (3.10)$$

where $\hat{u}_i^{(k)}$ is the next estimate of the variable, $\hat{u}_{temp}$ is the value of the component's response based on the values of all other components (as would normally be computed in the standard Gauss-Seidel solver), and $\Delta \hat{u}_i^{(k)}$ is the step (difference between the current and the previous update for the component). The relaxation factor is adapted on each iteration, as seen in Algorithm 2. Through dampening of the divergence, quicker convergence is expected [11]. XDSM diagram is given in Figure A.2.

The purely **monolithic Newton's solver** takes the residual of each component's governing equation and iteratively searches for the states that satisfy the coupled system residual equation. At each iteration, the solution is estimated by approximating the residual as a linear function (using the first two terms of the Taylor series, as shown in 3.3.2) and finding the step that drives the residual to zero. This requires solving a linear system at each iteration.

This approach directly accounts for all couplings but can be expensive or unstable if subsystems are strongly nonlinear. A slight variation implemented in OpenMDAO, the **full-space hierarchical Newton algorithm**, modifies it by inserting an inner stage before each outer Newton update: the subsystem nonlinear solvers are executed first to reduce their local residuals. These inner solves act like a nonlinear preconditioner, providing a better state on which the outer Newton step is applied. The outer step still solves the global Newton system using the full residual vector and Jacobian, ensuring that couplings between the subsystems are resolved. Thus, the hierarchical variant

⁷No algorithm and XDSM diagram is given here, as the next solver only introduces a small change, so the algorithm and diagram will be presented for these two algorithms together.

blends the robustness of local nonlinear solves with the quadratic convergence of a full-space (monolithic) Newton method. See Algorithm 3 and Figure A.3 for details.

In the remainder of the Chapter, the “Newton solver” and “Newton method” refer to the full-space hierarchical Newton algorithm variant.

The **Broyden solver** [200] modifies Newton’s method by using a finite-difference approximation of the Jacobian (or derivatives, in the 1-dimensional case). Instead of solving the linear system directly, it uses an approximated inverse of the Jacobian, which can make the method more efficient in certain cases. While the Broyden solver typically requires more iterations to converge, each iteration is computationally cheaper, as it involves fewer operations.⁸

3.3.2 Iterative solution example

The following example illustrates the problem formulation and computational steps of the Newton solver for *cycle 1*, as implemented in OpenMDAO, to support the brief introduction made in the previous section. It first outlines the theoretical basis of the Newton method in more detail, then defines the residual equations for the coupled system, and finally explains how the solver operates within the modular component framework of OpenMDAO.

The Newton method

Consider a nonlinear system of equations

$$R(u) = 0, \tag{3.11}$$

where $R \in \mathbb{R}^n$ is the vector of residuals and $u \in \mathbb{R}^n$ is the vector of unknown states.

The Newton method linearises the residual function about the current iterate $u^{(k)}$,

$$R(u^{(k)} + \Delta u) \approx R(u^{(k)}) + J(u^{(k)}) \Delta u, \tag{3.12}$$

where $J = \partial R / \partial u$ is the Jacobian matrix. Setting $R(u^{(k+1)}) = 0$ in the linear model

⁸The algorithm and diagram for the Broyden solver are the same as those for the Newton solver.

yields the update step Δu and the updated point $u^{(k+1)}$:

$$J(u^{(k)}) \Delta u = -R(u^{(k)}), \quad u^{(k+1)} = u^{(k)} + \Delta u. \quad (3.13)$$

Iterations continue until the residual norm $\|R(u^{(k)})\|$ is below a specified tolerance.

As outlined in Algorithm 3, the inner (hierarchical) solves of OpenMDAO reduce the local residuals of each subsystem to a specified tolerance before computing the system-level residual and Newton step.

System states and residuals in the coupled model

In the current model, the coupled problem is modelled using three components: *Equilibrium*, *Mooring*, and *Aero*. Each component contributes states and residual equations to the system assembled by OpenMDAO. The solver thus operates on a square system: the number of residual equations equals the number of unknown states.

The states correspond to the outputs of the components contained within the non-linear solver's scope:

- $\bar{\eta} \in \mathbb{R}^6$: platform displacement vector (*Equilibrium* component output).
- $\bar{F}_{moor} \in \mathbb{R}^6$: mooring force vector (*Mooring* component output).
- $\bar{F}_{aero} \in \mathbb{R}^6$: aerodynamic force vector (*Aero* component output).

Thus, the combined state vector is

$$u = \begin{bmatrix} \bar{\eta} \\ \bar{F}_{moor} \\ \bar{F}_{aero} \end{bmatrix}. \quad (3.14)$$

The *Equilibrium* component, modelled as an *ImplicitComponent* in OpenMDAO, enforces force balance between hydrostatic, gravitational, mooring, and aerodynamic loads on the floating platform. The residual is defined as

$$R_{\eta} = (C_h + C_g) \bar{\eta} - (\bar{F}_{aero} + \bar{F}_{moor}). \quad (3.15)$$

At equilibrium, $R_\eta = 0$, i.e. restoring forces balance external loads.

The Mooring component is modelled as an *ExplicitComponent* in OpenMDAO. Its residual is defined as

$$R_{\bar{F}_{moor}} = \bar{F}_{moor} - f_{mooring}(\bar{\eta}), \quad (3.16)$$

where $f_{mooring}(\bar{\eta})$ is the mooring load evaluated by MoorPy for a given displacement $\bar{\eta}$.

Similarly, the residual of the OpenMDAO's *ExplicitComponent* Aero is

$$R_{\bar{F}_{aero}} = \bar{F}_{aero} - f_{aero}(\bar{\eta}), \quad (3.17)$$

where $f_{aero}(\bar{\eta})$ is the aerodynamic load evaluated with CCBlade at mean displacement $\bar{\eta}$.

The complete system-level coupled residual vector is:

$$R(x) = \begin{bmatrix} R_{\bar{\eta}}(\bar{\eta}, \bar{F}_{moor}, \bar{F}_{aero}) \\ R_{\bar{F}_{moor}}(\bar{\eta}, \bar{F}_{moor}) \\ R_{\bar{F}_{aero}}(\bar{\eta}, \bar{F}_{aero}) \end{bmatrix} = \begin{bmatrix} (C_h + C_g) \bar{\eta} - (\bar{F}_{aero} + \bar{F}_{moor}) \\ \bar{F}_{moor} - f_{mooring}(\bar{\eta}) \\ \bar{F}_{aero} - f_{aero}(\bar{\eta}) \end{bmatrix}. \quad (3.18)$$

The corresponding Jacobian structure is block-partitioned as:

$$J = \frac{\partial R}{\partial [\bar{\eta}, \bar{F}_{moor}, \bar{F}_{aero}]} = \begin{bmatrix} C_h + C_g & -I & -I \\ -\frac{\partial f_{mooring}}{\partial \bar{\eta}} & I & 0 \\ -\frac{\partial f_{aero}}{\partial \bar{\eta}} & 0 & I \end{bmatrix}, \quad (3.19)$$

where I denotes the identity block of appropriate dimension. This Jacobian has contributions from the Equilibrium component (first row block) and from the Mooring and Aero components (lower rows). Each component provides only its local derivatives, and OpenMDAO conveniently assembles them into the global Jacobian.

Newton iteration process in OpenMDAO

As shown in Algorithm 3, at each nonlinear iteration, the following steps are executed:

1. State guess update: the solver holds a current guess for $(\bar{\eta}, \bar{F}_{moor}, \bar{F}_{aero})$.

2. Subsystems are solved using their local nonlinear solvers (implicit for *Equilibrium*, explicit for *Aero* and *Mooring*), using the most recent states from other subsystems. This reduces the local residuals and provides a better starting point for the global Newton update.
3. Residual assembly: the residual vector is assembled from each component's contributions:
 - Equilibrium residual: force imbalance at the current guess,
 - Mooring residual: mismatch between solver's guess of $\bar{F}_{moor}$ and the mooring model prediction
 - Aero residual: mismatch between solver's guess of $\bar{F}_{aero}$ and the aerodynamic model prediction
4. Jacobian assembly: each component provides its declared partial derivatives. OpenMDAO automatically transforms *ExplicitComponent* derivatives (outputs with respect to inputs) into residual derivatives and adds the identity terms on diagonal blocks.
5. Linear solve at the most recent state: Newton solves $J\Delta u = -R$ for the step Δu .
6. State update:

$$u^{(k+1)} = u^{(k)} + \Delta u. \quad (3.20)$$

7. Convergence check: iterations repeat until $\|R(u^{(k)})\|$ falls below tolerance.

3.4 Results and discussion

The convergence of cycle 1 does not depend on the solution of cycle 2 (and vice versa), allowing the best strategy for each cycle to be studied independently. For cycle 1, two tests were performed, each using one of the two solvers suitable for the system with an implicit component: Newton and Broyden. For cycle 2, five nonlinear solvers—NLBJ, NLBGS, NLBGS+A, Newton, and Broyden—were applied, tested, and compared. It is

important to note that the NLBJ and NLBGS solvers cannot be applied to cycle 1, as they require the equation of equilibrium (Equation 6.1) to be solved for η . This is not possible because the hydrostatic stiffness matrix is singular and has no inverse. The singularity arises from the lack of stiffness in the surge direction, which causes the first row of the matrix to be a zero vector.

The case considered here involves the OC3 Hywind spar platform [42] with the NREL 5MW reference wind turbine [43], subjected to rated wind speed (11.4 m s^{-1}) and a moderate sea state, with a significant wave height of 3.66 m and a peak period of 9.7 s. Subsection 3.4.1 compares the uncoupled and coupled models, while Subsection 3.4.2 presents the results from the solvers' benchmarking.

3.4.1 Influence of the introduction of the feedback loops

The coupled model, as introduced in Section 3.2.1 and Figure 3.2, includes two cyclic subsystems: cycle 1, coupling the mean position of the platform ($\bar{\eta}$) with the mooring and aerodynamic loads ($\bar{\mathbf{F}}_{\text{moor}}, \bar{\mathbf{F}}_{\text{aero}}$), and cycle 2, coupling the viscous damping with the significant deviation of the platform motion velocity ($\mathbf{B}_v, \sigma_u$). The comparative study presented here relates to the model with cycle 1 and cycle 2 present (coupled model), and the models with each cycle removed in turn, one at a time. The influence of eliminating coupling effects on the efficiency of the solution was studied by measuring the average simulation time over 100 runs.

The effect of cycle 1

As reported in Table 3.3, the introduction of coupling between the platform offset, mean aerodynamic load, and mean mooring loads increased the solution time by 16 %. While, in absolute terms, this increase is insignificant for a single simulation (on the order of seconds), it becomes more important when performing optimisations involving a large number of iterations or exploring an extensive design space for parametric analysis, in which case thousands of model evaluations might be needed. Additionally, the time cost of introducing the feedback loops depends heavily on the fidelity (and cost) of the models used within the cycles. The computations within the cycles were relatively

inexpensive, with the most time-consuming block (the BEM solver) executed after the first cycle converged.

Table 3.3: Performance of the uncoupled and coupled models - the effect of cycle 1.

		Coupled	Cycle 1 removed	Difference
CPU time	s	7.914	6.805	-16 %
Mean surge offset, $\bar{\eta}_x$	m	24.31	16.88	-31 %
Mean heave offset, $\bar{\eta}_z$	m	-0.40	-0.18	-56 %
Mean pitch inclination, $\bar{\eta}_\theta$	$^\circ$	4.88	4.35	-11 %
Aerodynamic load, $\bar{F}_{aero,x}$	N	6.85E+05	6.95E+05	1.5 %
Aerodynamic load, $\bar{F}_{aero,z}$	N	-5.99E+04	-6.01E+04	0.3 %
Mooring stiffness, $C_{moor,11}$	N m $^{-1}$	3.62E+04	4.12E+04	13.7 %
Mooring stiffness, $C_{moor,55}$	N m rad $^{-1}$	2.93E+08	3.16E+08	7.8 %
Max mooring line tension, T_{max}	N	1.27E+06	9.19E+05	-27.7 %
Nacelle acceleration, $\ddot{\eta}_{x,nacc}$	m s $^{-2}$	0.329	0.327	0.6 %

The coupling effects had a noticeable impact on the mean surge offset and pitch inclination angle: the uncoupled model underestimated these responses by 31 % and 11 %, respectively. This was primarily due to the model's inability to account for the nonlinearity of the mooring loads with respect to the platform displacement. In the uncoupled model, the offset is computed by equating the mean aerodynamic load with the product of the linear stiffness and displacement, evaluated at the undisplaced position of the platform. These non-conservative predictions should generally be avoided, as excessive platform displacements can lead to increased mooring line and power export cable loads. In fact, the uncoupled model underestimated the mean mooring tension by approximately 30 %. The uncoupled model also underestimated the heave offset by more than 50 %, as it did not account for the vertical component of the mean aerodynamic load on an inclined rotor and because it underestimated the vertical force from the mooring system.⁹

Importantly, the results were substantially improved by performing just one iteration of the Newton solver, reducing errors to below 10 % for all offsets and forces considered (except for the heave, where the improvement was smaller).

The significant amplitude of the nacelle acceleration was almost unaffected by the

⁹This large relative error corresponds to a small absolute error of 22 cm, small compared to the structural dimensions of the platform.

representation of the coupling effects. For a given FOWT design, it depends solely on the integral over a function of the motion RAO and the sea spectrum S_{ζ} (which is fixed), as shown in Equation 6.17. These factors were found to vary only slightly, except in the narrow region near the peak frequencies (Figure 3.4). The differences observed were attributed to variations in the aerodynamic damping (which is higher in the coupled system) and in the potential coefficients (which show slightly lower wave excitation and higher hydrodynamic added mass in the coupled system).




Figure 3.4: Comparison of the magnitudes of the surge and pitch RAO for the coupled and uncoupled models of dynamics. C1 - cycle 1, C2 - cycle 2.

Overall, the judgment of the trade-off between the gains in accuracy and the increase in computational effort is highly dependent on the fidelity of the models used both within and outside the cycle. In the case currently considered, the errors resulting from the uncoupled representation of the system were unacceptably high. However, a substantial improvement in the accuracy of the solution was achieved with a relatively small increase in computational effort.

The effect of cycle 2

The differences between the coupled and uncoupled solutions were marginal (below 0.5%) for the structure and load case studied, as seen in Figure 3.4. This was primarily due to the relatively low contribution of viscous damping to the total damping in the

sea state considered, as shown in Figure 3.5, which illustrates the effect of introducing subsequent sources of damping into the system. The introduction of this cycle increased




Figure 3.5: Surge RAO amplitude for different levels (sources) of damping: B_h - potential radiation damping, B_{aero} - aerodynamic damping, B_{add} - additional linear damping to calibrate with tank test, B_v - viscous damping.

the solution time by approximately 1 %. Considering the low additional cost, the cycle should be retained in the system. Although its influence on the results presented here was negligible, this may not be the case when different environmental conditions (e.g., higher sea states) and floating structures (particularly those with small-volume members) are considered.

3.4.2 Influence of the choice of the iterative solver

The previous test compared the results of the uncoupled and coupled models, with the coupled model using the Broyden and NLBGS+A solvers for its two cycles. In this test, the performance of the two iterative solvers for cycle 1 and the five iterative solvers for cycle 2, as defined in Section 3.3, is compared. The CPU performance counter (measured in fractional seconds) and the mean of the 100 results were recorded, based on the average execution time over 100 runs. For the study of cycle 1, the solver for cycle 2 was set to NLBGS+A, while for the study of cycle 2, the solver for cycle 1 was set to Broyden's method.

For this test, the frequency resolution was lowered to decrease the time cost of

running the BEM solver relative to the cost of performing the iterative solutions within the cycles. This adjustment allowed for a study of the performance of different solvers, with the results remaining relevant even if higher-fidelity solvers are applied within the cycles. In other words, instead of increasing the fidelity of the modules within the cycles, the fidelity of the modules outside the cycles was lowered to enable a fair comparison of the performance of the different solution strategies. The relative convergence tolerance (the change in the norm of residuals from one iteration to the next) was rigorously set at a very small value of 10^{-10} .¹⁰

As shown in Table 3.4, the choice of iterative solver influenced the solution time to a noticeable extent. Between the Newton and Broyden solvers applied to cycle 1, the latter was more time-efficient, reducing the solution time by -4% . For cycle 2, the NLBGS solver with Aitken relaxation correction was the most efficient, converging more than twice as fast as the same solver without the correction, and 60% faster than the NLBJ solver.

Table 3.4: Performance of nonlinear iterative solvers. C1 - cycle 1, C2 - cycle 2.

	Cycle 1		Cycle 2				
	Newton	Broyden	NLBGS	NLBGS+A	NLBJ	Newton	Broyden
Mean time (s)	8.254	7.914	20.156	7.914	17.699	15.231	9.936
(normalised)	1.00	0.96	1.00	0.39	0.88	0.75	0.49
NL iterations	8	10	5	4	25	4	6

The superior performance of the Broyden and NLBGS+A solvers is also evident in Figure 3.6, which shows the convergence history of the two cycles. All solvers, except for the NLBJ, achieved convergence in 10 or fewer iterations. For cycle 1, only minor differences were observed between the Newton and Broyden solvers, with the former converging three iterations earlier. This aligns with typical observations, where the Broyden solver generally requires more iterations, but each iteration is less time-expensive, as explained in Section 3.3.

¹⁰The relative convergence tolerance was chosen to ensure that the solver reaches a solution where the residuals are negligible, effectively guaranteeing that the iterative method has fully converged. At the cost of solution time, this is to avoid premature termination of iterations (with larger tolerances), which could result in a solution that does not satisfy the governing equation to the desired precision. By setting a very small tolerance, we ensure that the solver's final solution is close to the true root, avoiding any discrepancy due to insufficient convergence that would be propagated down the model.

Chapter 3. Multidisciplinary design analysis optimisation framework




Figure 3.6: Coupling variables convergence history for different solvers.

For cycle 2, the NLBGS+A and Newton solvers required only four iterations to converge. The NLBGS solver without the Aitken relaxation correction required just one more iteration, while the NLBJ solver took considerably more, requiring 25 iterations to converge (note that only the first 10 iterations are presented in the figure). Parallel computing was not enough to offset the disadvantage of the higher number of iterations required. As for the Broyden solver, it took two more iterations than the NLBGS+A and Newton solvers, but it converged the system significantly faster than the Newton solver.

As noted previously, a single iteration of all solvers, except for NLBJ, vastly improved the accuracy of the solution with only an almost negligible increase in the solution time.

It is also worth noting that the dominant source of nonlinearity driving the iterative convergence is the mooring system, as evidenced by the significant variation in $\bar{F}_{moor,x}$ across iterations, whereas the aerodynamic load $\bar{F}_{aero,x}$ remains nearly constant throughout. This is consistent with the expectation that the update of aerodynamic loads based on small platform pitch, whilst included for completeness, contributes only marginally to the overall convergence.

3.5 Conclusion

This section provides a summary of the chapter and outlines the specific contributions to knowledge made.

3.5.1 Summary

The study presented a coupled frequency-domain multidisciplinary model of dynamics. The framework of interconnected components (disciplinary analyses) was presented, following the principles of the multidisciplinary design analysis and optimisation (MDAO) approach. The most important coupling effects were modelled, including the interdependence of the mean offset of the platform and the mooring and mean aerodynamic loads, as well as the velocity of the movement of the platform and the viscous damping, which formed cyclic subsystems. The trade-off between the solution accuracy and efficiency for the coupled and uncoupled models was analysed, and a range of iterative solvers were benchmarked.

The coupling between the mean offset of the platform and the aerodynamic and mooring loads had a substantial influence on the values of the responses, converging at 31 % higher surge and 11 % higher pitch offsets, and at lower thrust and mooring stiffness. The uncoupled model underestimated the mooring tension by approximately 30 %. These non-conservative results demonstrate the criticality of the two-way coupling between the platform excursion and the mooring loads. The dynamic characteristics of the system's global motion response, such as the nacelle horizontal acceleration, were relatively unaffected by the introduction of the feedback loops. The numerical implementation of these coupling effects was achieved at a low increase in the total solution time (+16 %), due to the application of the efficient Broyden's method.

Regarding the coupling of the viscous damping with the statistics of the motion velocity, the discrepancies were marginal, mainly due to the low contribution of the viscous damping to the total damping and its minor effect on the response of the platform considered in the sea state considered. Although all solvers were very efficient, the best-performing Gauss-Seidel solver with Aitken relaxation massively outperformed

the other solvers tested.

Since the results of benchmarking of the different solvers are highly problem-specific, but relatively inexpensive, it is recommended to always perform such an investigation before attempting a more comprehensive study, such as design space screening or design optimisation, especially if time-intensive models are included within the cyclic subsystems.

Additionally, computing the potential coefficients for the platform in its equilibrium position, rather than the initial undisplaced position, was an improvement over state-of-the-art models; however, for the spar platform considered here and for small-amplitude heave and pitch displacements, this did not significantly affect the results.

3.5.2 Specific contributions

The following contributions to knowledge were made:

- Development of a modular FOWT multidisciplinary design analysis framework, with the coupled iterative solution for the mean platform offset, aero and mooring loads.
- Evaluation of the importance of converging the static equilibrium position under the mean aerodynamic and mooring loads and the additional computational cost.
- Evaluation of the importance of the viscous damping iterative solution for a spar FOWT and the additional computational cost.
- Comparison of the efficiency of alternative nonlinear iterative solvers and identification of the most efficient solution strategies for each coupled subsystem.

Chapter 4

Optimisation problem re-formulation

This chapter is partly based on the published paper:

P3 Patryniak K, Yeter B, Collu M, Coraddu A, Efficient design space exploration for spar Floating Offshore Wind Turbines: the trilemma of the ultimate, fatigue, and serviceability limit states (2024). *Journal of Physics Conference Series* 2767

4.1 Introduction

The design of floating offshore wind turbines (FOWTs) is carried out in accordance with established standards and recommended practices, in order to ensure that the structure can withstand relevant loads during installation, towing, operation, and damaged conditions, and operate safely throughout its lifetime [201]. The key standards include those issued by the International Electrotechnical Commission (IEC 61400-3-2 *Wind energy generation systems. Design requirements for floating offshore wind turbines* [202]), Det Norske Veritas (DNV-ST-0119 *Floating wind turbine structures* [201], DNV-RP-C205 *Environmental conditions and loads* [203], DNV-RP-0286 *Coupled analysis of floating wind turbines* [10]), and more recently, Lloyd’s Register (LR-RP-003-2024 *Recommended Practice for Floating Offshore Wind Turbine Support Structures* [204]).

In particular, the IEC 61400-3-2 standard provides a design basis by specifying 36 design load cases (DLCs) organised into 10 distinct design situations. These include power production, parked and redundancy scenarios, and checks under normal and fault conditions (among others). Each DLC is characterised by a unique combination of environmental and operational parameters, including wind, waves, current, water level, and system states. Within each DLC, a comprehensive set of specific loading conditions—defined as various combinations of these parameters—must be evaluated to ensure certification compliance for FOWTs. This amounts to a typical order of 10^4 load cases.

Under these conditions, a range of metrics must be evaluated to ensure the safety and reliability of the floating systems. In that regard, the offshore standards define a *limit state*: a condition beyond which a structure or structural component will no longer satisfy the design requirements. In particular, the ultimate limit state (ULS) corresponds to the maximum load-carrying resistance, fatigue limit state (FLS) corresponds to failure due to the effect of cyclic loading, accidental limit state (ALS) corresponds to survival in a damaged condition or in the presence of abnormal environmental conditions, and, finally, serviceability limit state (SLS) corresponds to project-defined criteria applicable to the intended use.

Within these categories, the most critical checks highlighted in DNV-RP-0286 [10] and DNV-ST-0119 [201] include:

- **ULS group:** loss of structural resistance (excessive yielding and buckling), brittle fracture, and permanent deformation (for all components, but in particular the tower base, tower to floater transition, mooring line and fairleads); loss of static equilibrium of the structure, tendons slack, ringing and springing.
- **SLS group:** maximum tilt and acceleration at tower top in operational and non-operational load cases, motions and vibrations affecting non-structural components, deformations that affect the efficient operation or watertightness, corrosion.
- **FLS group:** tower, floater, blades, and mooring line cumulative damage due to

repeated loads.

Although some metrics are not directly tied to the support structure design, they are influenced by the platform's motions, which are governed by it. Consequently, these metrics are viable candidates to serve as constraints in a FOWT support structure optimisation.

DNV proposes the design by the partial safety factor method, where design criteria are expressed in terms of (i) load and resistance factors (γ_F , γ_m) and (ii) characteristic values of load acting on the structure and resistance of the structure (F_k , R_k) :

$$\gamma_F F_k \leq R_k / \gamma_m \quad (4.1)$$

The above list of metrics and the proposed formulation provide a natural starting point for defining optimisation constraints. However, several challenges arise when extending a one-off design procedure into an optimisation problem. A flexible optimisation framework requires iterative design evaluations; while computationally expensive processes may be tolerable for a single design, they become impractical and unmanageable within an optimisation loop. In particular, some metrics require nonlinear time-domain analyses (e.g., maximum loads during transient events such as shutdown) or on high-fidelity methods such as finite element analysis (e.g., brittle fracture), both of which remain intractable at the system-level optimisation stage. Furthermore, the level of detail in the formulation of conceptual design optimisation constraints must match the information available at the conceptual stage: detailed structural features (e.g., internal scantlings) are typically not yet defined, with the emphasis instead on global scaling and preliminary capital cost estimation.

Similarly, the environmental conditions that define the DLCs in offshore standards provide an ideal reference set under which the design metrics (or constraints) should be evaluated. However, incorporating the full set of DLCs within an optimisation process quickly becomes infeasible. Conceptual-level studies often rely on a reduced set of DLCs, most commonly dominated by DLC 6.1 and DLC 1.6 in the FOWT design

literature [7, 205].¹ Even then, the number of required simulations can be substantial, as it grows with the factorial combination of environmental parameters (wind, wave, current, water level parameters) and the number of random seeds.

Therefore, both the set of design metrics and the environmental loading scenarios must be redefined and reduced to enable early-stage design optimisation. The remainder of this chapter addresses this need: Section 4.2 introduces a parametrisation of the FOWT support structure that balances level of detail with design flexibility; Section 4.3 defines objectives relevant to the conceptual design stage; and Section 4.4 presents a formulation of design constraints, grounded in offshore standards but adapted to ensure suitability for conceptual design and compatibility with efficient analysis modelling and optimisation approaches.

As discussed in Chapter 1, conceptual FOWT design optimisation may also benefit from metrics that are not explicitly required by offshore design standards but can provide additional insight into the global system behaviour. In this context, the instantaneous centre of rotation (ICR) is investigated in Chapter 5 as a potential unconventional design metric.

4.2 Design variables

As illustrated in Figure 4.1, the design variables encompass three subsystems: the tower, the floating platform, and the mooring system. The OC3 Hywind spar with a three-line catenary mooring system [149] serves as a reference for support structure design.

The degrees of freedom for the tower design include the tower base diameter and the ratio of the tower base diameter to wall thickness (d_{twr}/t_{twr}), hereafter referred to as the D/t ratio, also commonly also referred to as a slenderness or compactness² ratio. The tower base wall thickness is therefore derived from the diameter and the D/t ratio. It represents an effective thickness that accounts for additional structural

¹Such reduced sets of DLCs for the conceptual design stage have also been proposed in DNV-RP-0286 [10] and in LR-RP-003 [204].

²“Compact” in structural dynamics terms, meaning “thick-walled”.




Figure 4.1: Design parameters for the mooring line, floating platform and tower base section.

stiffness contributions—such as flanges, welds, bolts, and internal stiffeners—that are not explicitly modelled in this work. The tower top design remains unchanged from the original design to avoid modifying the tower–RNA interface (bedplate), and tower height is not altered as it is assumed to be one of the fundamental design requirements. Both diameter and wall thickness are assumed to vary linearly from the base to the top of the tower.

The floating support structure is parametrised through two variables: the draft and a vector of five diameters along the draft. This allows for some control over the water-plane area as well as the weight and underwater volume (buoyancy) distribution, which control the primary stability mechanism of the spar-type FOWT. The draft and shape of the lower part of the spar determine the ballast distribution, providing considerable flexibility in adjusting the system’s gravitational stiffness. Similar to tower, internal scantlings and secondary structures are not explicitly modelled. Instead, an effective steel density is applied to account for the additional structural mass associated with internal members, welds, and outfitting, following the approach adopted in the reference design of [149].

The mooring lines are parametrised through four parameters: chain nominal diameter, fairlead vertical location as a fraction of the draft, anchor radius, and mooring line shape parameter α_{moor} . The range of plausible line lengths depends on the

range of the fairlead and anchor positions, confined to the range between the shortest, diagonal line ($l_{min} = \sqrt{|z_F - z_A|^2 + |r_F - r_A|^2}$) and the longest, L-shaped line ($l_{max} = |z_F - z_A| + |r_F - r_A|$), as illustrated in Figure 4.2. The shape parameter is therefore formulated to allow all geometrically plausible shapes, as exemplified in Figure 4.3:³

$$\alpha_{moor} = \frac{l - l_{min}}{l_{max} - l_{min}} \quad (4.2)$$

The fairlead radius is dependent on the platform radius at the fairlead height. As in




Figure 4.2: Geometrically plausible mooring lines range.

Figure 4.3: Geometrically plausible mooring lines.

the original OC3 spar design, its value is set as the local value of the platform radius plus 0.5 m. The anchor vertical location is set by the water depth (location of the seabed).

The ranges of the variables (bounds in the optimisation problem) were selected to cover the range of designs that are likely to be considered in near-future projects, also considering the literature review presented in Section 2.2.1. They are summarised in Table 4.1 and detailed in the following:

- **Tower base diameter** ranges from the value equal to that of the tower top diameter (resulting in a non-tapered tower) to a value 40 % higher than the original, ensuring the design space is not overly restrictive.

³Note that even for $\alpha_{moor} = 0$, the mooring line will exhibit some curvature (or sag) due to chain elasticity. This effect, while small, is not zero. Accordingly, it is important to distinguish between the unstretched line length and the actual (stretched) line length. The formulations in this thesis are based on the theoretical unstretched line length, whereas Figure 4.3 presents the computed line shape.

- **The tower D/t ratio** of the original tower ranges from 203.7 at the top to 240.7 at the base. Here, values between 100 and 240 are investigated, corresponding to tower base designs that are never more compact (lower D/t) than the original top section but can be up to 100 % less compact (higher D/t). Such variations affect the system's centre of gravity, inertia, and tower vibration modes, as well as the structure strength and cost.
- **The platform draft** is capped at the reference value of 120 m, as deeper-drafted structures would be challenging to manufacture and deploy given port constraints. As shown in Figure 4.4, drafts below approximately 68 m would lead to a negative total pitch stiffness C_{55} (the sum of hydrostatic and gravitational contributions), causing the platform to lose initial stability. In addition, the draft must be sufficient to provide a restoring stiffness greater than that required to oppose the overturning moment from turbine rated thrust (about $720 \text{ MN m rad}^{-1}$ at a pitch angle of 5°), which is satisfied at a draft of roughly 100 m. For perspective, the lowest-draft spar-type FOWT currently installed, Hywind Scotland, has a draft of about 78 m, while typical designs range between 90 m and 120 m (e.g., OC3-Hywind reference design, Hywind Demo, Hywind Tampen, UMaine spar). These estimates depend on the specific platform, mooring system, the full FOWT configuration, and environmental characteristics; the practical draft range for optimisation in this work is set to 90 m–120 m.
- **Platform diameters** are limited to only slightly smaller than the original one up to 50 % larger than the largest diameter of the original design, which is well within the range usually considered for offshore steel structures. Larger diameters are of interest as they may skew the proportions of gravitational versus hydrostatic restoring moments, and also impact the effect of the mooring on the platform (through increased fairlead radius).
- **Fairlead vertical position** is expressed as a fraction of the platform draft to avoid incompatibility between the two variables. The range is set from the keel line ($z_F/T = 1.0$) up to just below the waterline ($z_F/T = 0.2$). Although fair-




Figure 4.4: Hydrostatics as a function of spar draft.

leads above the waterline could, in principle, improve the dynamic response, such configurations are not permitted here due to potential platform access constraints.

- **Anchor radius** The anchor radius is bounded to $\pm 30\%$ of the original design.
- **Mooring line shape parameter** is allowed to vary over its full generic range from 0 to 1, covering all geometrically plausible line shapes. Depending on the platform offsets under environmental loading, some of these shapes may not yield the desired static or dynamic response of the FOWT. Such solutions will be indirectly penalised during optimisation based on motion and load constraints.
- **Chain nominal diameter** range was decided considering the values of the two reference designs: the 5 MW design (OC3 spar [149]), as the smallest turbine that is still likely to be useful, and the 22 MW design (DTU and IEA Wind semisubmersible [206]), as the largest public floating wind turbine design up to date. The 5 MW design uses a 50 mm chain⁴, while the 22 MW design uses a 400 mm chain. However, based on the review of the technical specifications of offshore chain manufacturers (for example, Vicinay Marine [207]), currently,

⁴90 mm volume-equivalent diameter.

205 mm chain is the largest chain available. Therefore, the range 40 mm–200 mm is selected.

Table 4.1: Design variables and their bounds.

	Min	Max	Units
Tower			
Tower base diameter (d_{twr})	3.87	14.1	m
Tower base wall thickness to diameter (D/t) ratio	100	240	m
Floating platform			
Draft (T)	90	120	m
Diameters at $[0, 0.25, 0.5, 0.75, 1]T$ (d)	6.0	13.6	m
Mooring system			
Fairlead vertical position (z_F/T)	0.2	1.0	-
Anchor radius (r_A)	591.64	1098.76	m
Line shape parameter (α_{moor})	0	1	-
Nominal chain diameter (d_{moor})	0.04	0.2	m

4.3 Design objectives

The design objectives reflect the design goals and can vary significantly depending on the design stage, design priorities, and the capabilities of the analysis models. The conceptual design stage typically prioritises the development of a floating support structure that is feasible for the given offshore site and has a competitive CAPEX. Therefore, economic metrics are commonly adopted as objectives, leaving technical considerations to constraints.

Keeping the focus on the development of the efficient coupled dynamic model and its integration within a multidisciplinary optimisation framework, this work uses a simple mass-based metric as a proxy of the cost of the tower and floating platform:

$$f_{cost, str} = c_{steel}M_{steel} + c_{bal}M_{bal} \quad (4.3)$$

The cost factors account for the cost of raw material and manufacturing. The cost

of the mooring system is modelled as a product of the chain mass density m (mass per unit length of the chain), unstretched length l_{moor} , and an empirical factor:

$$f_{cost,moor} = c_{moor} \cdot m \cdot l_{moor} \quad (4.4)$$

The cost coefficient values are given in Table 4.2.

Table 4.2: Cost coefficient values for structure and mooring components.

Component	Cost coefficient	Value	Unit	Ref
Steel structure	c_{steel}	2.50	\$/kg	[6]
Ballast	c_{bal}	0	\$/kg	[6]
Mooring system	c_{moor}	2.585	\$/kg	[208]

4.4 Design constraints

This section focuses on the formulation of the design constraint functions; to keep it focused, only high-level formulations are provided, and the details of the models required to evaluate these formulations are presented separately in relevant chapters and appendices:

- Motions: Chapter 3,
- Modal analysis and flexible tower modelling: Appendix C,
- Structural loads: Chapter 6 and Appendix D.

The constraint metrics proposed for a concept-level support structure optimisation are described below, grouped by the limit state. For each group (SLS, ULS, FLS), the characteristic environmental conditions for evaluating the metrics are defined, as motivated by the DLCs' definition in IEC 61400-3-2, but significantly simplified to make the highly iterative optimisation process tractable.




Figure 4.5: Serviceability and ultimate limit state representative wind speed, significant wave height and peak period selection based on conditional probabilities.

4.4.1 SLS group

The serviceability metrics are often among the first checks performed during design feasibility studies. The SLS-group constraints are to be analysed under the rated wind speed condition, which is the design/target condition. Using metocean data from the Scottish sectoral marine plan site “NE8” [209], the sea state parameters were derived based on joint wind–wave statistics, as illustrated in Figure 4.5. The dataset was first filtered to include only records with wind speeds within a $\pm 1.0 \text{ m s}^{-1}$ interval around the rated wind speed. Within this subset, the significant wave height H_s was defined as the expected value of the empirical H_s distribution, and a range within ± 0.25 m of this value was selected. From within this range of records, the peak period T_p was then taken as the expected value of the corresponding empirical T_p distribution. The resulting environmental conditions are summarised in Table 4.3.

Although more rigorous approaches to environmental condition selection exist (for example, the IFORM multivariate contours recommended in IEC 61400-3-2 [202]), these methods remain an active field of research. Since the focus of this work is not on environmental contour derivation, simplified but statistically consistent conditional methods were adopted to provide representative inputs for the analyses.

The SLS metrics considered are outlined below, along with their respective limiting values. The values are based on the consensus established in the literature (see

Table 4.3: Environmental conditions. V_s - mean hub-height wind speed, TI - turbulence intensity, H_s - significant wave height, T_p - wave peak period, γ - wave shape parameter, θ_w - mean wave direction, V_c - current speed. ULS_1 and ULS_2 refer to conditions defined in Section 4.4.2.

LS	V_s	TI	α	H_s	T_p	γ	θ_w	V_c
	(m s^{-1})	(%)	(-)	(m)	(s)	(-)	($^\circ$)	(m s^{-1})
SLS, ULS_1	11.4	14	0.2	1.87	7.48	1.0	0	0
ULS_2	24.6	14	0.2	5.91	9.86	2.96	0	0
FLS	3.0–25.0	14	0.2	1.06–5.91	7.39–9.86	1.0–2.96	0	0

Section 2.2.1 Table 2.5).

Motions constraints

In this section, the “maximum” response quantities are approximated using the mean response plus three standard deviations within a given short-term environmental condition, consistent with the underlying linear Gaussian response assumptions adopted in the frequency-domain framework. These quantities therefore represent statistical estimates of extreme response levels rather than true time-domain maxima or long-term design extremes associated with a specified return period.

1. Estimated maximum offset: excessive rigid-body offsets may cause damage to the power export cable. The constraint is formulated as the norm of the displacement vector composed of surge, sway, and heave displacements, expressed as a fraction of water depth d :⁵⁶

$$g_{\text{offset}} = \|\eta_{\text{max},1:3}\| \leq 0.15d \quad (4.5)$$

2. Estimated maximum tilt: excessive tilt may lead to inefficient operation of the RNA components or loss of static stability. The constraint is formulated as the

⁵This is a general formulation; however, only the surge and heave translational modes are modelled here.

⁶Note that since the cable attachment point is not in general located at the origin where the translational degrees of freedom are defined, rotational modes would also contribute to the offset experienced at that point. However, this coupling is not accounted for in the present formulation of the offset constraint, the limiting value of which is set conservatively.

norm of the displacement vector composed of roll and pitch:⁷

$$g_{\text{tilt}} = \|\eta_{\text{max},4:5}\| \leq 12.5^\circ \quad (4.6)$$

3. Estimated maximum resultant acceleration at the hub level: uncontrolled vibration or acceleration may lead to inefficient operation of the non-structural components: yaw motor and brake, yaw, pitch and main bearings, fluid levels, and cooling fluid movement, also driving the structural loads in the structural components (for example, contributing to the fluctuating bending moment at the tower base):

$$g_{\text{acc}} = \|\ddot{\eta}_{1:3}\| \leq 0.3g \quad (4.7)$$

Natural frequencies constraints

The risk of large loads due to the resonance of the exciting loads with the eigenfrequencies of a FOWT makes the analysis of the system frequencies necessary. The offshore standards recommend obtaining the coupled system's natural frequencies (both damped and undamped) from either a frequency-domain coupled model or a Fast Fourier Transform of time-domain simulations. The modes most relevant to the design of the support structure include: first and second global tower bending, floating substructure surge, sway, heave, roll, pitch, and yaw, as well as the first tower torsion mode. The standards also recommend developing a Campbell diagram to map where the structural natural frequencies coincide with the expected excitation frequencies. A safety margin of at least $\pm 5\%$ to crossings of structural eigenfrequencies with 1P (rotor frequency), 3P (blade passing), and 6P (twice the blade passing) excitation should be kept [10].

In this work, a simplified and more time-efficient modal analysis is adapted to establish the system's natural frequencies, focusing instead on a more comprehensive and robust formulation of the natural frequency constraints for concept design optimisation purposes. The floating platform and rotor-nacelle-assembly are treated as a rigid bod-

⁷This is a general formulation; however, only the pitch rotational mode is modelled here. In principle, yaw should also be constrained; however, for the case of the spar-type FOWT with a symmetric mooring system and unidirectional environmental load, this is insignificant.

ies, while the tower is modelled as a beam with distribution of structural properties and geometry along its height.⁸ The system's rigid body in-plane modes⁹ and the first tower fore-aft mode are compared against the environmental loading (wind and first-order wave loads) and rotor characteristic (1P and 3P) frequencies. For each frequency regime, a penalty function is introduced that reflects the risk of natural frequencies being close to the excitation ranges. Within the optimisation, these penalty functions provide a quantitative measure of constraint violation, guiding the design away from resonance-prone regions.

The wave-frequency penalty function is derived by first establishing a probability distribution of the wave peak period using metocean data [209]. A histogram is generated (panel (1) in Figure 4.6), capturing the peak periods and their corresponding probabilities of occurrence across the whole dataset (no wind speed or wave height conditioning). For each of 22 bins in the histogram, the mean of the peak period and the mean of the significant wave height are identified, and the JONSWAP power spectrum is computed (panel (2) in Figure 4.6).¹⁰ To incorporate the likelihood of each peak period, the power spectra are then scaled by the probabilities from the peak period histogram, resulting in weighted spectra that account for both the power content and the probability of occurrence (panel (3) in Figure 4.6). The resulting probability-weighted spectra are summed (panel (4) of Figure 4.6) and normalised to the range 0–1 to produce a continuous and smooth penalty function (panel (5) in Figure 4.6), representing the likelihood and severity of overlapping with the wave frequency range.

The wind range frequency penalty function is constructed following the same procedure, but using the Kaimal wind spectrum scaled with the probability of occurrence of a given mean wind speed bin (mean value from the bin of mean wind speeds), as illustrated in Figure 4.7. 22 bins are used.

The rotor frequency penalty was formulated as a smooth and continuous function,

⁸Tower flexibility is only considered in the evaluation of tower natural frequencies. While the influence of platform dynamics on tower modes is considered, the effect of tower flexibility on platform dynamics and structural loads is not. This is further discussed in the relevant chapters.

⁹This is consistent with all models across the thesis only concerning the in-plane motions and loads, however, the method is general and could be applied to investigate arbitrary modes.

¹⁰The wave peak-shape parameter was set based on the peak period and significant wave height, as recommended in the IEC 61400-3 Annex B [210].




Figure 4.6: The process of deriving the natural frequency penalty for the first-order wave excitation region.

suitable for robust use in numerical optimisation. The process is illustrated in Figure 4.8 and explained below.

Panel 1 in Figure 4.8: the turbine control curve specifies the rotor speed Ω (in revolutions per minute) as a function of wind speed U . A long-term wind speed time series can be mapped through this curve, giving a rotor speed value at each time step. Conversion from rpm to Hz, and multiplication by the harmonic index (1 or 3), yields two datasets of operating frequencies: the 1P and 3P excitation frequencies (only 1P is considered in the figure for illustrative purposes).

Panel 2: Histogramming the 1P (or 3P) dataset produces an empirical probability density function $p(f)$, which describes how often the rotor operates at each frequency. This distribution tends to be flat and relatively low in the low-frequency range, with a concentration at higher frequencies. This reflects the dominance of the near-rated wind speed in the metocean dataset and the corresponding operation of the rotor at its rated speed (12.1 rpm, 0.2017 Hz), as well as flattening of the control curve in the




Figure 4.7: The process of deriving the natural frequency penalty for the wind excitation region.

above-rated conditions.

Panels 3 and 4: the penalty is designed to quantify how risky it is for a system's natural frequency f_{nat} to lie near the rotor's excitation frequencies. A large penalty should result if f_{nat} coincides with or lies close to frequencies where the rotor spends much of its operating time, and a small penalty otherwise. It should therefore account for two factors simultaneously:

- proximity to excitation frequencies, and
- probability of occurrence of the excitation frequencies.

Hence, the penalty at a given f_{nat} is defined as:

$$P(f_{nat}) = \int p(f) K(|f_{nat} - f|) df, \quad (4.8)$$

where $p(f)$ is the empirical density and $K(d)$ is a kernel function that decreases




Figure 4.8: 1P frequency range penalty derivation process.

smoothly with distance $q = |f_{\text{nat}} - f|$. The kernel is chosen as a Gaussian:

$$K(q) = \exp\left(-\frac{q^2}{2\sigma^2}\right), \quad (4.9)$$

which decays smoothly, matches the intuition that resonance influence diminishes with frequency separation, and provides the differentiability required for gradient-based optimisation. The parameter σ (kernel width¹¹) has the interpretation of a frequency margin: larger values imply that frequencies further away are still considered risky. Here, it is set to 0.009 Hz.

Panel 3 shows several example kernels, while panel 4 shows the convolution terms:

$$p(f)K(|f_{\text{nat}} - f|) \quad (4.10)$$

illustrating the convolution mechanism.¹² Each of these curves represents the local contributions that are integrated to yield the penalty value at that f_{nat} .

Note that the kernel centred at 0.23 Hz (yellow curve in panel 3) results in a zero convolution term (no visible shade in panel 4) due to the lack of overlap with the empirical 1P distribution (panel 2). On the contrary, the kernel centred at 0.16 Hz fully overlaps with 1P distribution, resulting in a large convolution component (dark green shade in panel 4).

To avoid artificial truncation of the penalty function at the operating frequency limits, the convolution is evaluated on a slightly extended frequency grid (padded at both ends by 3.5 kernel widths). This ensures that the penalty decays smoothly to zero outside the occupied rotor frequency band, rather than cutting off abruptly at the edges.

Panel 5: the convolution is evaluated over a grid of candidate natural frequencies spanning the 1P or 3P range, resulting in a smooth penalty curve. Each point on this curve corresponds to the sum of products $p(f)K(|f_{\text{nat}} - f|)$ across all frequencies f , i.e.

¹¹ σ is the standard deviation of the Gaussian kernel, more precisely.

¹²The convolution term $p(f)K(|f_{\text{nat}} - f|)$ represents the weighted contribution of each rotor frequency f to the penalty at a candidate natural frequency f_{nat} . It is defined as the product of the empirical rotor frequency distribution $p(f)$ and a kernel function K centred at f_{nat} , and its integral (the area under the curve) gives the penalty value at f_{nat}




Figure 4.9: Comparison between the traditional (fixed region with $\pm 5\%$ margins [10]) and the new proposed resonance constraints for the rotor characteristic frequencies.

the area under the convolution term shown in panel 4.

Panel 6: the penalty curve is normalised so that its maximum equals one, ensuring all penalty values lie in the range $[0, 1]$. In this way, values near unity indicate a high resonance risk, while values near zero indicate “safe” frequencies.

In current state-of-the-art approaches, the optimiser is strictly prevented from selecting designs with natural frequencies within the 1P and 3P ranges, regardless of how likely the rotor is to operate at those frequencies. The proposed probability-aware formulation instead applies smooth, differentiable penalties weighted by the likelihood of excitation. As illustrated in Figure 4.9, this still protects the 1P and 3P regions, but in a soft manner that permits solutions within these bands when the probability of excitation is low. If a design solution slightly violates the resonance constraint but is highly beneficial in terms of other design metrics, an optimiser might still consider it a good candidate solution. Such a treatment not only provides greater flexibility in the design space—potentially leading to technically and economically superior solutions—but also yields smooth gradients that are more suitable for gradient-based optimisation methods, as will be discussed further in Chapter 7.

An important consequence of this approach is that the artificial split of designs into soft (below 1P), soft–stiff (between 1P and 3P), and stiff–stiff (above 3P) categories can be relaxed. Under the traditional formulation, the hard barrier created by the 3P range

meant that optimisers were effectively confined to whichever class the initial design belonged. Hegseth et al. [51] showed that local cost minima exist in both the soft–stiff and stiff–stiff regions for the first tower bending mode, but the 3P barrier prevented transitions between them. By introducing smooth probability-weighted penalties, the proposed formulation removes this restriction and enables exploration across categories, offering greater potential for global design improvements.

Note that this is a quasi-static approach that assumes the rotor speed follows the control curve exactly. In practice, however, platform motions, limits on blade pitch rate, and other dynamic effects cause the rotor to deviate from the prescribed curve for short periods, resulting in under- and overshoots. These deviations are not explicitly captured in this approach, although the safety margin introduced through the kernel width parameter σ partly compensates for them.

When evaluating resonance avoidance constraints, the system natural frequencies are queried against each precomputed penalty landscape (wind, wave, 1P and 3P), as illustrated in Figure 4.10. Separate constraints are formulated for each natural frequency and for each excitation range, allowing the optimiser to identify specific design changes that shift natural frequencies away from risky regions. In the present problem, this results in 16 inequality constraints of the form:¹³

$$P(f_{\text{nat}}) \leq 0. \quad (4.11)$$

which steer the solution outside penalised frequency ranges gently, rather than imposing rigid cut-off bands as has been common in the literature.

¹³Technically, the penalty function cannot exceed zero. The constraint is expressed as an inequality rather than an equality to remain consistent with the optimisation problem formulation and solution approach described in Chapter 7.




Figure 4.10: Wind, wave, 1P, and 3P range penalty functions with the reference system natural frequencies (dashed lines: surge, heave, pitch, and tower 1st fore-aft modes).

Anchor vertical pull

Additionally, the vertical load on the anchors is constrained to account for the limited capacity of drag-embedded anchors to resist vertical loading:

$$F_{z,moor} \leq 30N. \quad (4.12)$$

A relatively small limiting value of 30 kN is imposed, corresponding to approximately 3.5 % of the pretension in the original design. While offshore standards and recommended practices typically require zero vertical load at the anchor, a small positive allowance is introduced here to improve optimiser performance by avoiding constraint enforcement at the boundary of a flat region in the design space. Such limited vertical loading is not considered detrimental according to multiple sources in the literature, depending on the anchor technology and soil conditions [211].

4.4.2 ULS group

The ULS-group constraints are analysed under two conditions: ULS_1 , the rated wind speed corresponding to the maximum thrust condition, and ULS_2 , the wind speed just below cut-out, corresponding to the maximum wave height condition. The sea state parameters were selected based on joint wind–wave metocean statistics, as described in

Section 4.4.1 and illustrated in Figure 4.5. While multivariate extreme value analysis is typically employed to identify ULS cases, in this work a simplified yet systematic approach is adopted.

Tower base yielding: The von Mises stress at the tower base, $\sigma_{\text{VM},\text{Sd}}$, is compared against the material yield stress, f_y :

$$\sigma_{\text{VM},\text{Sd}}\gamma_F \leq \frac{f_y}{\gamma_M}, \quad (4.13)$$

accounting for γ_F : the load factor, and γ_M : the material partial safety factor. The stress $\sigma_{\text{VM},\text{Sd}}$ accounts for axial load from the RNA and tower self-weight, bending moment induced by aerodynamic loading on the rotor and tower, inertia effects due to acceleration and inclination of the mass items above the tower base, and in-plane shear from aerodynamic loading. A detailed description of the yielding check is provided in Appendix D.3.

Tower shell buckling. Because the tower shell is thin compared to its diameter, it can fail prematurely by buckling at stresses lower than the material yield, under axial compression (resulting from the RNA and tower weight, as well as aerodynamic loads) and bending (resulting from aerodynamic loads and inertia). Therefore, the tower is checked for longitudinal shell buckling under the combined action of axial force and global bending, assuming a circular cylindrical shell without longitudinal stiffeners in accordance with DNV-RP-C202 [212]. Circumferential membrane stress from external pressure and in-plane shear stress from torsion/shear are neglected because they are small relative to the longitudinal compressive stress. The utilisation is evaluated at the most compressed location by comparing the design longitudinal stress to the design buckling strength:

$$\sigma_{x,\text{Sd}}^{\max} = \sigma_{a,\text{Sd}} + \sigma_{m,\text{Sd}} \leq f_{kd}, \quad f_{kd} = \frac{f_{ks}}{\gamma_M}. \quad (4.14)$$

Here $\sigma_{a,\text{Sd}}$ and $\sigma_{m,\text{Sd}}$ are the design axial-membrane and bending stresses, respectively, and f_{ks} is the characteristic shell-buckling strength for a circular cylindrical shell in longitudinal compression obtained per DNV-RP-C202 from the elastic buckling stress

and material yield strength, accounting for geometry (D/t , L/D), boundary conditions, and fabrication tolerances (imperfection class). Further details of the shell buckling assessment are given in Appendix D.1.

Global column buckling. In addition to local shell instability, the tower is also verified against global column buckling, i.e. flexural instability of the entire tower treated as a cantilever column. This mode is usually not governing for modern offshore wind turbine towers, which are relatively thick-walled, but it is checked for completeness, since the tower diameter and thickness are the design variables.

The axial load N_{Sd} and bending moment M_{Sd} at the tower base are compared against the design axial and bending resistances, respectively:

$$\frac{N_{Sd}}{N_{Rd}} + \frac{|M_{Sd}|}{M_{Rd}} \leq 1.0, \quad (4.15)$$

with

$$N_{Rd} = \frac{\chi A f_y}{\gamma_M}, \quad M_{Rd} = \frac{W f_y}{\gamma_M},$$

where A is the gross cross-sectional area, W is the section modulus, χ is the reduction factor from the Eurocode buckling curve, and γ_M is the material partial safety factor. Further details of the column buckling assessment are given in Appendix D.2.

Mooring line tension: The constraint is formulated following the recommendation of DNV-OS-E301 [213], based on the mean and dynamic tension contributions (T_{mean} , T_{dyn}) with corrosion correction. The tension in a mooring line at the fairlead location (in the first chain link), augmented by the safety factor (SF), should not exceed the minimum breaking strength corrected for corrosion:

$$(\gamma_{mean} T_{mean} + \gamma_{dyn} T_{dyn}) / SF \leq MBL(d_{corr} / d_{new}) \quad (4.16)$$

A corrosion rate of 0.2 mm/year over half of the design life (12.5 years) is assumed.

4.4.3 FLS group

The environmental conditions for evaluating the FLS-group constraints were selected from metocean data covering wind speeds from cut-in to cut-out, divided into 2 m s^{-1} bins (the maximum allowed by DNV-OS-E301 [213] and IEC 61400-3 [214]). As illustrated in Figure 4.11, the midpoint of each bin was taken as the representative wind speed. For each wind speed bin, the representative H_s was defined as the expected value of the conditional distribution $H_s | V_s$. To ensure consistency, the dataset was further filtered within $\pm 0.25\text{m}$ of this H_s , and the expected value of T_p was then computed. This sequential conditioning yields self-consistent sea state representations (V_s, H_s, T_p) for use in fatigue analysis.





Figure 4.11: Fatigue limit state representative wind speed, significant wave height and peak period selection based on conditional probabilities.

2 m s^{-1} bins correspond to 11 bins in the operational range. The effect of the limited number of bins was assessed by computing the total fatigue damage using progressively finer binning, based on the model defined in the next chapter. As shown in Figure 4.12, reducing the bin width below 2 m s^{-1} has a relatively small influence on the predicted fatigue damage.¹⁴

Although the directionality of the environmental loading is, in general, important when assessing fatigue, this model assumes a single wind and waves direction of 0° and applies a fatigue reduction factor to approximately account for the directionality of the waves (distribution of the mean wave direction), as proposed in [215]. The metocean

¹⁴At the conceptual design level, this level of discretisation is considered sufficient; however, detailed design would warrant a more comprehensive investigation of appropriate fatigue conditions binning.

Figure 4.12: Fatigue bin width sensitivity. 2 m s^{-1} bin marked in black.

data is used to obtain the probability of each mean wave direction. Then, the correction factor is obtained according to Equation 4.17 as demonstrated in Table 4.4.

$$R_{\mu} = \sum_{i=1}^I \left[p(\mu_i) |\cos \mu_i|^3 \right] \quad (4.17)$$

Table 4.4: Directionality correction factor calculation.

μ_i ($^{\circ}$)	p (%)	R_i
15	13.2 %	0.1277
45	5.3 %	0.0375
75	3.3 %	0.0085
105	7.9 %	0.0205
135	8.5 %	0.0598
165	10.8 %	0.1045
195	9.0 %	0.0868
225	5.6 %	0.0396
255	6.6 %	0.0170
285	9.2 %	0.0239
315	8.6 %	0.0610
345	12.0 %	0.1155
$R_{\mu} = \sum_{i=1}^I R_i$		0.7023

The fatigue analysis will be further detailed in Chapter 6. In the following, the FLS

metrics adapted as constraints are briefly introduced.

Tower base fatigue damage The fatigue damage at the tower base section is computed by summing up the contributions from all environmental conditions bins, under the assumption of linear cumulative damage (Palmgren-Miner rule):

$$D_{twr} = T_d \sum_{i=1}^N D_{i,twr} p_i R_\mu \quad (4.18)$$

T_d is the design life of the FOWT in hours, D_i and p_i are the contribution to 1 h fatigue from the i -th bin and the bin's probability, respectively, R is the directionality reduction factor.

The fatigue damage $D_{i,twr}$ is computed in accordance with DNV-RP-C203 [26] using the hotspot stress range, which is defined as the product of the nominal stress—arising from inertial and aerodynamic loads—and the stress concentration factor, accounting for geometric discontinuities such as the slope between sections of varying diameters. The corresponding constraint is formulated as:

$$D_{twr} \leq 1/DF F_{twr} \quad (4.19)$$

where DFF is the Design Fatigue Factor (DFF) from DNV-OS-C101 [216], reflecting the probability and consequence of failure. For the tower base section, a value of $DF F_{twr} = 1$ is used, considering that it is an external structure accessible for inspection, located above the splash zone (more than 5 m above the waterline), with a 5-year inspection interval conducted on-site afloat.

Mooring line fatigue damage The tension-tension fatigue damage (excluding the out-of-plane bending effects) of the most loaded mooring line at the first link of the chain (connected to the fairlead) is computed as:

$$D_{moor} = T_d \sum_{i=1}^N D_{i,moor} p_i R \quad (4.20)$$

The stress required for the fatigue assessment is computed accounting for a corrosion rate of 0.2 mm/year over half of the design life (12.5 years). Additionally, following

the recommendation of DNV-OS-E301 Section 6.5.4 [213], an additional load factor $\gamma_F = 3$ is applied to account for uncertainties in the fatigue analysis, assuming that the mooring system is inspected at a suitable offshore or onshore facility. Therefore, the constraint is formulated as follows:

$$D_{moor} \leq 1/\gamma_F \quad (4.21)$$

4.4.4 ALS group

ALS corresponds to survival conditions during extreme loading events and addresses system safety following local damage caused by accidental loads [108]. Load analysis under ALS conditions examines transient phenomena resulting from events such as emergency stops or mooring line failures, which cannot be accurately modelled using fast engineering tools, such as the frequency-domain model developed in this work. Consequently, ALS is typically not considered in early concept or feasibility studies and will not be further addressed in this thesis.

4.5 Conclusion

This section provides a summary of the chapter and outlines the specific contributions to knowledge made.

4.5.1 Summary

This chapter presents a re-formulation of the FOWT support-structure design optimisation problem. The formulation is grounded in offshore standards and recommended practices, but adapted for conceptual-level design, where design information is limited and fast evaluation methods are required. The design variables define key features of three subsystems—the tower, floating platform, and mooring system—chosen to balance flexibility with a tractable number of degrees of freedom. The objectives represent cost estimates for these subsystems. The constraints include metrics that serve as prox-

ies for the safety checks recommended in offshore standards. The complete optimisation problem formulation is summarised in Table 4.5.

Table 4.5: Summary of design formulation: variables, objectives, and constraints.

Category	Subsystem	Quantity	Symbol / Expression
Design variables	Tower	Base diameter	d_{twr}
		D/t ratio	$\alpha_{\text{twr}} = d_{\text{twr}}/t_{\text{twr}}$
	Platform	Draft	T
		Diameters	$d(z)$ at $z/T \in \{0, 0.25, 0.5, 0.75, 1\}$
	Mooring	Fairlead vertical pos.	$\alpha_F = z_F/T$
		Anchor radius	r_A
		Line shape parameter	α_{moor}
		Nominal chain diam.	d_{moor}
Objectives	Support str.	Cost proxy	$f_{\text{cost, str}} = c_{\text{steel}}M_{\text{steel}} + c_{\text{bal}}M_{\text{bal}}$
	Mooring	Chain cost	$f_{\text{cost, moor}} = c_{\text{moor}} m l$
	Mooring	Total length	$f_{l, \text{moor}} = 3l$
Constraints	SLS	Max offset	$\ \eta_{\text{max}, 1:3}\ \leq 0.15d$
		Max tilt	$\ \eta_{\text{max}, 4:6}\ \leq 12.5^\circ$
		Hub accel.	$\ \ddot{\eta}_{1:3}\ \leq 0.3g$
		Nat. freq. penalty	$P(f_{\text{nat}}) \leq 0$
		Anchor vert. pull	$F_{z, \text{moor}} \leq 30kN$
	ULS	Tower yielding	$\sigma_{\text{VM}, \text{Sd}}/f_{yd} \leq 1.0$
		Tower shell buckling	$\sigma_{a, \text{Sd}} + \sigma_{m, \text{Sd}} \leq f_{ks}/\gamma_m$
		Tower column buckling	$\frac{N_{\text{Sd}}}{N_{\text{Rd}}} + \frac{ M_{\text{Sd}} }{M_{\text{Rd}}} \leq 1.0$
		Mooring strength	$\frac{\gamma_{\text{mean}}T_{\text{mean}} + \gamma_{\text{dyn}}T_{\text{dyn}}}{\text{SF}} \leq \text{MBL}\left(\frac{d_{\text{corr}}}{d_{\text{new}}}\right)$
	FLS	Tower base fatigue	$D_{\text{twr}} \leq 1/\text{DFF}_{\text{twr}}$
		Mooring fatigue	$D_{\text{moor}} \leq 1/(\text{DFF}_{\text{moor}} \gamma_F)$

The problem formulation developed here will form the basis for the dynamic model development (Chapter 6) and optimisation case study (Chapter 7).

4.5.2 Specific contributions

The following contributions to knowledge were made:

Chapter 4. Optimisation problem re-formulation

- Formulation of the optimisation problem grounded in offshore standards and suitable for the conceptual design stage.
- New parametrisation of the catenary mooring line, including a shape parameter which allows to simultaneously vary the mooring line length and fairlead and anchor positions, ensuring the geometrical plausibility of the system (compatibility of the three variables).
- New formulation of the resonance avoidance constraints, based on both the probability and power content of the exciting frequencies (wind, waves, 1P, 3P).

Chapter 5

Unconventional metric: instantaneous centre of rotation

This chapter is based on two published papers:

- P5** Patryniak K, Collu M, Coraddu A, Rigid body dynamic response of a floating off-shore wind turbine to waves: Identification of the instantaneous centre of rotation through analytical and numerical analyses (2023). *Renewable Energy* 218
- P6** Patryniak K, Collu M, Jonkman J, Hall M, Barter G, Zalkind D, Coraddu A, Investigation into the instantaneous centre of rotation for enhanced design of floating offshore wind turbines (2025). *Wind Energy Science* 10

5.1 Introduction

As discussed in Chapter 2.2.1, compact system-level metrics capable of representing the complex dynamic behaviour of FOWTs may be valuable for design interpretation and optimisation. Considering the instantaneous centre of rotation (ICR) as one such potential metric, this chapter investigates the characteristics, behaviour, and potential design relevance.

Specifically, addressing the gap identified in Section 2.2.2, this chapter proposes methods to identify the ICR of FOWT and its statistics (Section 5.2) through time-

domain, dynamic numerical approach.¹ The dynamic analysis in regular and irregular waves is carried out to understand the motion patterns and ICR sensitivity to wave period and height. Then, the ICR of design variations under the complex loading conditions typical of operational FOWTs is further investigated to provide understanding supporting the design process. The initial step involves examining the influence of more complex environmental loading conditions (Section 5.4) to enhance understanding of ICR behaviour in realistic operational scenarios. For ICR insights to be useful in the design process, two conditions must be met:

- The ICR must be correlated with FOWT responses that are critical for design improvement.
- A set of design variables that can be adjusted to “control” the ICR statistics must be identified to enable these improvements.

Accordingly, Section 5.5 explores how specific design decisions affect ICR behaviour, followed by an investigation of the relationship between the ICR and key FOWT responses (Section 5.6). This sensitivity analysis helps identify key design features that can be prioritised during the FOWT optimisation process. This leads to the definition of a practical design scenario, which will be presented later in Chapter 6.5.

This Chapter focuses on the analysis performed for the spar-type FOWT. However, corresponding analysis of the semisubmersible-type FOWT is available in previously published work [217].

5.2 Methodology

This section outlines the methods used to derive and study the ICR of FOWTs. The dynamic model, environmental conditions, and coordinate systems used to simulate the motion of the FOWT are introduced, followed by the statistical methods required to analyse the ICR. Finally, the general methodology for ICR analysis is outlined.

¹Note that the models used in this chapter are not the ones developed and used in previous chapters. Previous chapters focused on developing frequency-domain and hybrid domain models, while this section uses a well-established time-domain model to study the ICR.

5.2.1 ICR derivation

When exposed to collinear unidirectional wind, wave, and current loads, FOWTs undergo general plane motion, with translational and rotational characteristics that depend on the external loads and the structure’s design. In particular, the centre of the pitch rotation depends on the amplitudes, frequencies, phases, and distributions of the external dynamic forces, and therefore its determination is nontrivial.

A point (not necessarily within the body) of zero velocity at a particular instant in time (i.e. ICR) can be found by analysing the velocity vectors at two arbitrarily chosen points on the body. By definition, the ICR appears to be instantaneously at rest, all other points in the body present pure rotation about this point; i.e. they follow circular paths around the ICR [110]. Therefore, the ICR can be found at the intersection of the lines perpendicular to the velocity vectors at the two points (A and B in Figure 5.1), assuming a rigid body. In the case when the velocity vectors are parallel, of different magnitudes, and do not cross and the line joining points A and B is perpendicular to the velocities, the ICR can be obtained through a direct proportion² [110], as illustrated in the middle subfigure of Figure 5.1. Finally, in the special case of pure translation (right subfigure of Figure 5.1), the perpendiculars to the velocity vectors do not cross, and so the ICR tends to infinity.

Note that the location of the ICR changes as the body moves or loading conditions change, in terms of both the body-fixed and global inertial coordinate systems, and therefore, it only exists at a particular instant in time (hence, “instantaneous”).

The vectors required for this calculation can be found by measuring the velocities in A and B , which are obtained from the time-domain numerical simulations (detailed in Section 5.2.2).

5.2.2 Dynamic simulations

All simulations are performed with OpenFAST [133], a state-of-the-art, extensively validated nonlinear aero–hydro–servo–elastic coupled time-domain model of dynam-

²Direct proportion is a relation between two quantities where the ratio of the two is equal to a constant value.




Figure 5.1: Construction of the ICR using velocity vectors at two points on the floating body. Left: general case, misaligned nodes, middle: special case, aligned nodes, right: special case, pure translation.

ics [218]. Equations of motion of OpenFAST are nonlinear and are expressed in a modular generalised state-space form with module-level states and module-to-module coupling handled through input–output constraints [219]. The degrees of freedom include six rigid-body motion modes; although possible in OpenFAST, the elastic dynamics are not modelled, as the analysis relies on a rigid-body assumption. The aerodynamic solution relies on the blade-element momentum theory with Beddoes–Leishman unsteady blade airfoil aerodynamics model, tip and hub loss corrections. Turbulent wind excitation was included according to the IEC 61400-1 Ed. 3 recommended spectrum. The first-order potential hydrodynamics are modelled through frequency-to-time-domain transforms based on the potential coefficients obtained from the boundary element method (BEM) code WAMIT [220], as available in [221]. The geometry of the spar platform was not varied within the study; therefore, one set of coefficients is sufficient. Viscous loads are computed from Morison’s theory, with the hydrodynamic coefficients listed in Table 5.1. Second-order hydrodynamics are not modelled. The sum-frequency

Table 5.1: Hydrodynamic model parameters.

Parameter	Unit	OC3 spar
Maximum Morison element size	m	0.5
Maximum BEM panel size	m	1
Added mass coefficient	-	-
Transverse viscous-drag coefficient	-	0.6

loads act at frequencies exceeding the dominant natural frequencies of the spar system, so although important from the flexible structure response point of view, they do not excite significant rigid-body motions [147]. The difference-frequency loads for

the spar are, in general, at least an order of magnitude smaller than the first-order hydrodynamic loads [147, 150], and also small as compared to aerodynamic thrust, in the order of 1 % of the mean thrust [150]. The nonlinear mooring loads are computed using either a quasi-static or a lumped-mass model.³ The initial investigation uses a quasi-static model MAP++ [222, 223]. The code calculates steady-state forces on the mooring lines' segments, including the effects of the seabed contact (kinetic friction force calculated as the product of a kinetic friction coefficient and the seabed contact normal force). The catenary line profile and effective forces are computed based on the distributed line mass, strain, and elasticity, for a line suspended in static equilibrium (hence *quasi-static*). Forces arising from the inertia, viscous drag, internal damping, bending and torsion are all ignored. The latter investigation of the ICR sensitivities is performed using a lumped-mass mooring line model, MoorDyn [224]. The model includes the effects of axial stiffness and damping, weight and buoyancy, hydrodynamic viscous forces, and vertical spring-damper forces from contact with the seabed.

5.2.3 Environmental conditions

The fundamental study presented in Section 5.3 considers a range of regular and irregular wave periods and heights outlined in Table 5.2. The selected wave periods were

Table 5.2: Case study definition - wave conditions.

Waves type (-)	(Peak) period (s)	(Significant) wave height (m)
Regular	10.47–125.66	2.0, 4.0, 8.0
Irregular	10.47, 31.25, 125.66	2.0, 4.0, 8.0

chosen to span the low-frequency response regimes most relevant to rigid-body motion of FOWTs, including conditions near the surge, pitch, and heave natural periods, in order to isolate and interpret the fundamental mechanisms governing ICR behaviour.

The more detailed study presented in Section 5.4 investigates the ICR under increasingly realistic environmental conditions, including shorter irregular waves and combined

³The latter investigation was an extension of the initial research and used a generally more accurate model.

environmental loading. Instead of relying on the predefined design load cases (e.g. DLC 1.6, which involves an operational turbine in normal wind and severe sea state conditions, and DLC 6.1, which considers a parked turbine experiencing extreme wind and sea conditions, as outlined in the IEC 61400-3-2 standard [225]), this study selectively applies different loading components. This approach aims to provide deeper insights into the ICR behaviour related to specific sources of loading. The wind conditions are characterised by the mean hub-height wind speed, profile (shear), and turbulence intensity, and the wave conditions (sea states) are characterised by the (significant) wave height and (peak) period. Wind and wave are co-directional and aligned with the positive x -axis. The wave peak-shape parameter is set based on the peak period and significant wave height, as recommended in the IEC 61400-3 Annex B [210].

Sections 5.5 – 5.7, on the other hand, investigate ICR over a range of realistic environmental conditions. Each design variant is simulated in 10 environmental conditions, listed in Table 5.3, representing binned metocean data obtained for the Scottish sectoral marine plan for the NE8 site [209] (these bins account for 98.4% of cumulative probability). The responses presented in this chapter refer to the average responses weighted by the probability of each of the 10 conditions. Each simulation is run for 1 hour (excluding the initial transient phase) with a single random seed (the same seed in all simulations).

Table 5.3: Environmental conditions binning. Wind shear fixed at 0.2 and turbulence intensity at 14%.

Condition	1	2	3	4	5	6	7	8	9	10
V_s (m s^{-1})	4.1	6.3	8.5	10.7	12.9	15.1	17.3	19.5	21.7	23.9
H_s (m)	1.09	1.23	1.44	1.75	2.17	2.68	3.27	3.95	4.71	5.42
T_p (s)	8.68	8.54	8.14	7.65	7.51	7.63	8.04	8.58	9.16	9.58
Probability	0.087	0.213	0.266	0.230	0.134	0.052	0.013	0.002	0.001	0.001

5.2.4 Floating system and coordinate system

The results presented here relate to the case of the OC3 Hywind spar platform with a catenary mooring system, as defined in [226], and the NREL 5 MW reference wind

turbine, as defined in [43].⁴ The details of the reference design are given in Tables 5.4–5.5 and Figure 5.2.

Table 5.4: OC3 Hywind spar FOWT original design characteristics. CG – centre of gravity; CB – centre of buoyancy.

Property	Unit	OC3 spar
Platform, including ballast:		
Draft	m	120
Diameter	m	6.5–9.7 (tapered)
Offset column diameter	m	-
Base column diameter	m	-
Mass	kg	$7.47 \cdot 10^6$
Vertical CG	m	-89.92
Pitch inertia about CG	kg m ²	$4.23 \cdot 10^9$
Hydrostatics:		
Vertical CB	m	-61.17
Metacentric height	m	28.7
Stiffness coefficient:		
Heave	N m ⁻¹	$3.33 \cdot 10^5$
Pitch	N m rad ⁻¹	$-4.99 \cdot 10^9$
Infinite-frequency added mass:		
Surge–surge	kg	$7.569 \cdot 10^6$
Heave–heave	kg	$2.35 \cdot 10^5$
Pitch–pitch	kg m ²	$3.7 \cdot 10^{10}$
Surge–pitch	kg m	$-4.71 \cdot 10^8$
Mooring system:		
Fairlead height	m	-70
Fairlead radius	m	5.2
Water depth (anchor position)	m	-320
Anchor radius	m	853.87
Unstretched line length	m	902.20
Chain diameter (volume-equivalent)	m	0.09
Pretension	kN	910
Natural frequency:		
Surge	Hz	0.008
Heave	Hz	0.032
Pitch	Hz	0.034

This study employs two distinct coordinate systems. The global inertial coordinate system, illustrated in Fig. 5.2, offers a reference frame for examining how the ICR responds to varying environmental loading conditions and design parameters. By remaining unaffected by the structural displacements, it provides a clear view of how

⁴OC3: offshore code comparison collaboration; OC4: offshore code comparison collaboration continuation

Table 5.5: NREL 5 MW reference rotor characteristics.

Property	Unit	Value
RNA mass	kg	350000
Hub height	m	90.0
Cut-in, rated, cut-out wind speed	m s^{-1}	3, 11.4, 25
Rated wind speed thrust	kN	805.11
Tower:		
Elevation to tower base	m	10.0
Elevation to tower top	m	87.6
Tower mass	kg	249 718
Tower vertical CG	m	43.4

external forces influence the ICR. It is centred at the intersection of the mean waterline and the undisplaced tower axis (centroid of the undisplaced initial waterplane area). Conversely, the body-fixed coordinate system is used for design-focused evaluations, where the goal is to position the ICR relative to specific locations on the floating structure. This system initially coincides with the global inertial coordinate system but moves (both translating and rotating) with the floating body. Throughout the chapter, the results are primarily presented in the global inertial frame, except where the local frame is used to support the interpretation of observed design trends (where appropriate).




Figure 5.2: FOWT design variables and coordinate system.

5.2.5 ICR statistics

Since the ICR can exhibit asymptotic behaviour (even during a well-behaved single-frequency sinusoidal motion - as will be shown later) and is very irregular under stochastic environmental loading conditions, analysing time-domain signals directly can be difficult and not readily applicable to the design context. Therefore, a statistical analysis approach is followed instead.

The ICR is evaluated over the final 10 pitch periods to characterise representative behaviour while excluding the initial transient response and maintaining consistent coverage of the periodic response. Since the coupled FOWT motions contain multiple frequency components, the selected window does not fully isolate all dynamic contributions and therefore represents an approximation. Nevertheless, the adopted window was found sufficient to capture the dominant ICR trends and distributions discussed in this study.

The distribution characteristics of the in-plane coordinates of the ICR (type, mean, mode, expected value, standard deviation, skewness and kurtosis) are investigated. The mode (position of the peak of the distribution) gives a useful indication of the most often encountered values. However, in the cases where the ICR distributions are multimodal, the expected value is a more useful statistic. It is defined as the value that is the most likely result of the next repeated trial of a statistical experiment [227] and uses the probabilities of possible outcomes as weights in the process of finding the weighted average of the data in a data set. Here, it is adapted to find the average ICR coordinates weighted by the frequency of occurrence in the time signal recorded over 10 pitch periods, sampled at 0.0125s. Additionally, time histories and histograms of the ICR and rigid motions are presented to complete the picture.

The one-sample Kolmogorov–Smirnov (K–S) test with hypothesis testing [228] is performed to establish whether the ICR follows a normal distribution. This test compares data sampled from a fitted normal distribution with the samples from the original data and outputs two metrics: the K–S statistic, which measures the maximum discrepancy between the empirical cumulative distribution functions of the two samples, and the p-value, which indicates the probability of observing the test results under

the null hypothesis. The null hypothesis is that the two samples come from the same distribution. A high p-value above the significance level of 0.05 suggests insufficient evidence to state that the two samples come from different distributions. Therefore, a low K–S value and a p-value above 0.05 are needed to confirm that the data closely follow a normal distribution.

To analyse the relationships between design variables, the ICR, and FOWT responses, the Spearman rank correlation coefficient⁵ is used, which quantifies the statistical dependence between two variables. As illustrated in Table 5.6, for each pair of variables (for example, the standard deviation of the tower base bending moment and the standard deviation of z_{ICR}), their values from different simulations are listed and ordered (the smallest value in a list gets rank 1, the second-smallest value gets rank 2, etc.). The Spearman correlation coefficient ρ_s is computed based on the sum of the differences in the ranks of the paired values:

$$\rho_s = 1 - \frac{6 \sum_{i=1}^n d_{rank,i}^2}{n(n^2 - 1)} \quad (5.1)$$

$d_{rank,i}$ is the difference between the ranks for sample (simulation) i , and n is the number of observations (simulations).

Table 5.6: Spearman rank correlation calculation example.

Sample	M_y std (kN m)	Rank	z_{ICR} std (m)	Rank	$d_{rank,i}$
1	2.65E+03	5	120.55	4	1
2	2.66E+03	4	119.29	5	-1
3	2.87E+03	3	126.24	2	1
4	3.21E+03	2	122.14	3	-1
5	3.72E+03	1	149.58	1	0

Spearman’s correlation assesses the strength of a monotonic relationship, making no assumptions about the underlying data distribution. A correlation close to +1 or −1 indicates a strong monotonic relationship:

- $\rho_s \approx 1$: as one variable increases, the other tends to increase.

⁵The rank of a value is its position within a sorted list of data points, ordered from smallest to largest.

- $\rho_s \approx -1$: as one variable increases, the other tends to decrease.
- $\rho_s \approx 0$: there is little to no monotonic relationship between the variables.

Since Spearman's correlation is based on ranks rather than raw values, it is robust to nonlinearity and less sensitive to outliers than Pearson's correlation. However, it does not imply causation, meaning that a strong correlation does not indicate that changes in one variable directly cause changes in another.

5.2.6 Analysis approach

The remainder of this chapter will use the numerical method introduced in Section 5.2.1 to investigate the ICR of FOWTs in synthetic and real operational environmental loading conditions. As outlined in Figure 5.3, the analysis starts with a series of time-domain analyses carried out in regular and irregular waves to study the time-dependent behaviour of the ICR and to establish its sensitivity to wave frequency and height in a simple scenario of wave-only loading following the approach derived in Section 5.2.1. Building on that, the next, more involved and applied analysis stage, depicted in the bottom part of Figure 5.3, aims to i) understand the behaviour of the ICR in a complex, realistic environment, and ii) demonstrate how knowledge of ICR behaviour can be leveraged to improve FOWT designs.

The ICR can serve as a useful metric in the FOWT design process, provided that

1. the FOWT design can be modified in ways that influence the ICR statistics, and
2. the ICR statistics significantly affect the responses of interest (i.e. the performance of the FOWT).

Therefore, to facilitate the practical application of new knowledge about the ICR of FOWTs, this study examines the relationships among design variables, ICR statistics, and response statistics. In particular, design variables significantly correlated with the ICR statistics, and ICR statistics significantly correlated with the responses of interest are investigated, as illustrated in the middle blocks of Figure 5.3.




Figure 5.3: The procedure for investigating the ICR of FOWT, including an initial theoretical study (P5) and a more detailed study aimed at practical application (P6).

The study begins with one-at-a-time perturbations of multiple design variables, which are features of the floating platform and mooring system. This approach allows for the examination of correlations with ICR statistics (the mean and standard deviation of horizontal and vertical ICR components). Although more comprehensive experimental designs, such as full-factorial or Latin hypercube, offer deeper insights into design space behaviour, perturbing one variable at a time while keeping all other variables at their original design values allows for a clear presentation and interpretation of the results and significantly reduces the design space size, keeping the computational cost tractable.

Once the relationships between the design variables and ICR are established, the most impactful parameters are identified for down selection. These parameters are then subjected to a more thorough sensitivity analysis using a full-factorial approach. This analysis explores the relationships between all combinations of the selected parameters and the responses of interest (i.e. design performance metrics).

At this stage, highly correlated pairs of design variables and ICR statistics are further refined to include only those that also show strong correlations with the responses of interest. This systematic procedure helps identify design directions that may lead to




Figure 5.4: Wave energy spectra for 9 sea conditions (JONSWAP).

significant improvements in the performance of the FOWTs through changes to ICR statistics, supporting the definition of the design case studies (last block of Figure 5.3).

5.3 ICR in waves

The method described in 5.2.1 was applied to study the time history and statistics of the ICR coordinates during the motion of the moored platform in regular and irregular waves. The simulations were carried out for a range of wave periods and heights outlined in Table 5.2 in Section 5.2.3. The JONSWAP frequency spectrum was used to generate irregular waves, as recommended by the IEC 61400-3 Annex B [214] and shown in Figure 5.4. The results of the two series of tests will be presented in turn.

5.3.1 Regular waves

The simulations were carried out in regular waves of three different heights (2 m, 4 m, and 8 m), at periods ranging from 10.47 s to 125.66 s. Figure 5.5 presents the time history of the horizontal and vertical coordinates of the ICR (x_{ICR} and z_{ICR}) during the platform's motion in 2m waves of frequency equal to the natural pitch frequency (0.2 rad s^{-1}), plotted together with the pitch motion signal. The transient part of the signal is excluded, and the last 10 pitch periods are shown. Both x_{ICR} and z_{ICR} present periodic behaviour with tangent-like and parabolic shapes, which are the consequence of the pitch motion changing direction twice every cycle. At instances when the pitch




Figure 5.5: ICR time history in pitch natural frequency regular waves motion for 10 pitch periods.

velocity is zero (i.e., pitch displacement reaches maximum or minimum), the ICR becomes infinity. The physical interpretation can be made by examining a single pitch period, as displayed in Figure 5.6. Four phases of the pitch motion can be distinguished, each with characteristic ICR behaviour:

- Phase 1: as the pitch angle increases from zero to maximum, the x -coordinate of the ICR decreases from 19.5 m to minus infinity. The infinite x_{ICR} at the pitch motion peak corresponds to an instantaneous lack of rotation or a pure translation. As mentioned before, this is to be expected as, at this point, the pitch motion changes direction, hence the platform momentarily becomes static. z_{ICR} presents a similar behaviour, except, it increases to plus infinity at the end of the phase. A discontinuity is seen in both coordinates shortly after half of the phase - as will be shown later in this section, this is the effect of the mooring lines restraining the motion of the platform.
- Phase 2: as the pitch angle decreases from maximum to zero, the x -coordinate of the ICR decreases from plus infinity to a finite value of 22.2 m. z_{ICR} follows an




Figure 5.6: ICR time history in pitch natural frequency regular waves motion for 1 pitch period, derived from Figure 5.5.

opposite trend, increasing from minus infinity to a negative value of -61.9 m.

- Phase 3: as the platform pitches towards the opposite direction, x_{ICR} repeats the same pattern as in phase 1. z_{ICR} increases to plus infinity, with a discontinuity shortly after the half of the phase.
- Phase 4: as the platform returns to its initial position, x_{ICR} decreases back to 19.5 m, and z_{ICR} increases back to -61.9 m. Note the anti-symmetry of the x_{ICR} plot.

Both coordinates span infinite ranges. Although the above analysis aids the understanding of the dynamics of the platform pitch motion, it does not allow for the identification of a point of minimum additional translation due to the rotation during general motion, which is of practical interest. For this reason, the statistics of x_{ICR} and z_{ICR} are examined next, as listed in Table 5.7.

The positive value of mode of x_{ICR} indicates motion asymmetry - the ICR concentrates 27.98 m aft of the centre of flotation (i.e., the origin of the assumed coordinate system). The consequence of this is that the side of the floater facing the incoming

Table 5.7: Statistics of the ICR components based on the last 10 pitch periods, for the pitch natural frequency.

Statistic	x_{ICR} (m)	z_{ICR} (m)
Mean	30.73	-101.35
Mode	27.98	-102.38
Expected value	20.61	-62.27
Standard deviation	3019.99	264.24

waves experiences a higher amplitude of motion. The mode of z_{ICR} (-102.38 m) is at a depth not coinciding with any of the commonly considered characteristic points (i.e., the centre of flotation, buoyancy, gravity and the fairlead attachment point). The large vertical distance from the most often encountered centre of rotation from the hub impacts the velocities and accelerations encountered at the hub level. Both coordinates have a high standard deviation, indicating that the values are spread out over a wide range. While x_{ICR} distribution is almost perfectly symmetric around the mean, z_{ICR} is skewed away from the mean. This is a consequence of the fact the z_{ICR} of FOWT in regular waves follows a bi-modal distribution, as illustrated in Figure 5.7, where the space centroe of ICR (or locus: a curve formed by all ICR points) is plotted together with the probability distributions.

The space centroe of ICR consists of four curves, with the points concentrated near the centre of the plot. This shows that, at this frequency, the rotational behaviour of a FOWT is significantly more complex than that of FPSO vessels, the centroides of which have been reported to be straight lines in multiple references, including [121] (refer to the literature review in Section 2.2.2). Hence, unlike for FPSO, the derivation of an analytical model for the ICR centroe would be very challenging.

Lastly, the dependence of the ICR coordinates on the frequency and height of the incoming regular waves was studied.

Wave frequency effect A clear dependence of both ICR coordinates on the frequency of the waves was observed (Figure 5.8). $E(x_{ICR})$ peaks at a frequency slightly higher than the surge natural frequency, following a downward trend thereafter. An opposite tendency is observed for $E(z_{ICR})$. In the low-frequency limit, x_{ICR} and z_{ICR}




Figure 5.7: ICR space centred with probability distributions of the x and z coordinates. The coordinates of the ICR were z-score normalised (their distribution has a mean of 0 and a standard deviation of 1).

tend to infinity, which corresponds to a pure translation. The value of ICR increasing with the frequency decreasing is consistent with the behaviour observed previously in [117], where the large ICR values were associated with the condition of the long waves, whereby a floating body with a long natural period follows the profile of the sea surface, which corresponds to the motion dominated by translation, with minimal rotation (pitch). In the high-frequency limit, both $E(x_{ICR})$ and $E(z_{ICR})$ tend to zero (i.e., the centre of flotation). At the heave natural frequency (0.201 rad/s), $E(x_{ICR})$ peaks slightly due to the resonance between the rigid DOF and the oscillatory wave load, achieving values of 28.32 m to 47.66 m, depending on the wave amplitude.⁶ Note that these values do not coincide with either the centre of flotation/gravity or the mooring attachment point. Divergence from the values obtained by the analytical approach (*initial centre of rotation*) is also observed (recall the result of the analytical approach: $x_{ICR} = z_{ICR} = 0.0$ m). See Figures G.1 and G.2 in Appendix G for a comparison of the time series of ICR coordinates at different frequencies.

⁶This situation is similar to that illustrated in the right panel of Figure 5.1—for a given horizontal velocity, increasing the vertical velocity leads to the x-coordinate of the ICR shifting away from the body. This will be better explained in Section 5.4.2.




Figure 5.8: The expected value of ICR as a function of frequency for different wave heights.

Wave height effect As can be seen in Figure 5.8, the expected values of both coordinates present significant sensitivity to the wave height at low frequencies and almost no sensitivity to the wave height at the frequencies higher than the pitch natural frequency. This characteristic behaviour at the low frequencies can, among other factors, be attributed to:

- Mooring loads, nonlinear in platform offset: a catenary mooring system provides the station-keeping ability in the function of the weight of the suspended line length: when the platform is displaced away from the equilibrium position, a portion of the upwind line resting on the seabed is lifted, increasing the length of the free-hanging chain and changing the angle of the lines, which nonlinearly affects the tension in the lines and the load on the floating platform [36].⁷ The changes in mooring stiffness due to the platform offsets impact the low-frequency surge motions [32, 33].
- Viscous drag, nonlinear in wave amplitude: the viscous drag depends on the square of the relative velocity, i.e., the sum of the water particle velocity and the platform motion velocity [159], dependent on the wave height. The response of the FOWT near the resonant frequency (surge natural frequency) is highly dependent on the viscous loading on the platform [229, 230].

⁷Note that this is true for the lines opposite to the direction of the platform movement; the opposite is true for the other lines.

Therefore, to explain the low-frequency discrepancies, the expected value of ICR has been recalculated modifying the underlying physical model to exclude the two nonlinear effects. The nonlinear mooring loads were eliminated by replacing the quasi-static mooring model with a linearised one (fixed stiffness matrix obtained at the undisplaced platform position). The nonlinear viscous loads were eliminated by setting the viscous drag coefficient within OpenFAST's strip theory approach to zero.

As presented in Figures 5.9 and 5.10, both nonlinear effects (mooring and viscous loads) had a significant effect on the sensitivity of ICR to the wave height. Replacing the quasi-static mooring model with a linear one led to the reduction in the discrepancies in the z_{ICR} at all but the lowest two frequencies considered. Removing the viscous drag from the hydrodynamic model resulted in a significant reduction in the discrepancies in x_{ICR} at all frequencies. The two effects combined in such a way that the sensitivity of the ICR to wave height was greatly lowered (Figures 5.9c and 5.10c).




Figure 5.9: Impact of nonlinear effects on $E(x_{ICR})$.

The part of the mooring load linear in platform displacement was also found to contribute to the observed relationship between the ICR and wave height. Figure 5.11 presents $E(x_{ICR})$ and $E(z_{ICR})$ as functions of wave frequency for different wave heights, obtained for an unmoored structure.

A significant reduction in discrepancies at low frequencies can be seen, as compared to the standard systems' results (Figure 5.8). Additionally, when the mooring load




Figure 5.10: Impact of nonlinear effects on $E(z_{ICR})$.




Figure 5.11: Expected value of ICR as a function of frequency for different wave heights H , obtained for an unmoored system. The surge, heave, and pitch natural frequencies are marked with vertical lines for reference.

was removed from the model, both coordinates of ICR at 0.05 rad s^{-1} changed their signs, making the $x_{ICR}(\omega)$ and $z_{ICR}(\omega)$ curves monotonic. One more time, a significant impact of the mooring system on the dynamic behaviour of FOWT in low-frequency waves has been revealed.⁸

Since the ICR is computed based on velocity vectors, its sensitivity to the wave height primarily stems from the sensitivity of the motion velocity to the wave height.

⁸An effect not directly discussed in the above is that of the surge natural frequency "vanishing" in the absence of the mooring system; due to no stiffness in the surge direction, the only oscillation in the surge direction is due to the pitch rotation, and one can no longer speak about the resonance with the waves in the low-frequency range, which might contribute to the observed reduction in the ICR sensitivity to wave amplitude.

Indeed, Figure 5.12b shows the highly nonlinear relationship between the wave height and the platform motion velocity (measured at the system's origin), but only in the low-frequency case — this is not seen in the higher frequency (10 s) case (Figure 5.12a).




Figure 5.12: Comparison of the surge and heave displacement and velocity time histories for two different wave frequencies.

As noticed before, the complex motion behaviour at the low frequency (i.e., in long waves) is partly due to the effect the moorings have on the system. In the coupled 6 DOF motion, due to large surge offsets in near-resonant conditions, the mooring loads lead to the distortion of the heave and pitch motion patterns, as can be observed in Figure 5.13a.

The nature of this distortion depends on whether the mooring load is linear or nonlinear in platform displacement, and can be understood through the concept of *signal clipping* borrowed from signal processing [231, 232].

Consider the pitch displacement of a freely floating (unmoored) platform in steady-




Figure 5.13: The time history and FFT plots of the pitch motion obtained with different dynamic models.

state response to regular waves: since the hydrodynamic forcing is at a single frequency ω and the system is linear in the absence of moorings, the pitch motion is purely sinusoidal,

$$\theta(t) = \Theta \sin(\omega t), \quad (5.2)$$

where Θ is the pitch amplitude. The linear mooring stiffness restrains the platform's surge motion, which feeds back into pitch through the geometric surge–pitch coupling. Crucially, a linear (constant) stiffness matrix is symmetric: it provides equal restoring force for positive and negative platform displacements. This symmetry means that the distortion it imposes on the pitch signal satisfies $f(-\theta) = -f(\theta)$, i.e., it is an odd function of displacement. When an odd-symmetric distortion is applied to a sinusoidal signal $\sin(\omega t)$, the Fourier series of the resulting waveform contains only *odd harmonics*

of the excitation frequency:

$$\theta(t) \approx a_1 \sin(\omega t) + a_3 \sin(3\omega t) + a_5 \sin(5\omega t) + \dots \quad (5.3)$$

This is the so-called *hard clipping*: the linear mooring stiffness flattens the peaks of the pitch signal symmetrically, which resembles clipping in an audio amplifier. The result is visible both in the time history (Figure 5.13a, middle row) and in the FFT (Figure 5.13b, middle row), where harmonics at odd multiples of the wave frequency appear.

The nonlinear (catenary) mooring, by contrast, introduces an *asymmetric* restoring force: when the platform surges in one direction, the upwind mooring line tightens nonlinearly while the downwind line slackens, so the stiffness is greater in one direction than the other at large amplitudes. This asymmetry means $f(-\theta) \neq -f(\theta)$, i.e., the distortion is no longer an odd function, and the Fourier series of the pitch response acquires *both odd and even harmonics* (as well as a non-zero mean offset a_0):

$$\theta(t) \approx a_0 + a_1 \sin(\omega t) + a_2 \sin(2\omega t) + a_3 \sin(3\omega t) + \dots \quad (5.4)$$

This is *soft clipping*: the positive and negative peaks of the pitch signal are rounded off differently, as visible in the bottom rows of Figures 5.13a and 5.13b.

The strength of both clipping effects increases with wave height, since larger waves drive larger surge offsets and therefore stronger mooring-induced distortion.

A complementary source of nonlinearity is the viscous (Morison) drag, which is quadratic in the relative velocity between the water particles and the platform (u_{rel}):

$$F_D \propto |u_{\text{rel}}| u_{\text{rel}}. \quad (5.5)$$

For a sinusoidal relative velocity $u_{\text{rel}} = U \sin(\omega t)$, the drag force is an odd-symmetric function of u_{rel} , and its Fourier expansion contains only odd harmonics of ω . In the unmoored case, the platform moves freely with the long waves, keeping the relative velocity small and drag largely insignificant. Once the linear mooring stiffness is in-

troduced, the platform is restrained and no longer tracks the wave motion freely; the relative velocity increases, activating the quadratic drag as a meaningful source of odd harmonics. This effect is further amplified at higher wave heights, since larger wave particle velocities increase u_{rel} and therefore the drag force, consistent with the wave-height sensitivity seen in Figure 5.13b.

Both effects — mooring-induced clipping and quadratic drag — share the same root cause: the large surge motion at low frequencies, which grows with wave height. Although the linear mooring stiffness is itself a linear addition to the system, the equations of motion solved by OpenFAST are nonlinear: by restraining surge while pitch remains primarily governed by hydrostatic restoring forces, the linear mooring stiffness creates a phase imbalance between the two DOFs that is absent in the unmoored case. It is this mooring-imposed phase difference that activates the geometric surge–pitch coupling and introduces harmonic distortion — confirmed by the fact that disabling the surge–pitch coupling⁹ suppresses the even harmonics, as depicted in Figure H.1 in Appendix H.

5.3.2 Irregular waves

The study of the ICR of FOWT subjected to regular wave loading provided a good basis for the investigation of its behaviour in irregular waves, being one step closer to a real operational scenario which is of practical interest.

As shown previously in Figures 5.5–5.6, in regular waves (a purely theoretical condition), the ICR coordinates exhibit a distinctive, regular time history pattern. When the pitch displacement reaches its minimum or maximum, the pitch velocity is zero (pure translation), and the ICR tends toward infinity. At other points, the ICR concentrates around finite locations. The phases before and after each pitch extremum cause the ICR to cluster around unique points, as illustrated in the histograms of Figure 5.7, resulting in a bi-modal distribution. In a more complex, stochastic environment, this regular pattern is not observed (Figure 5.14). However, the ICR coordinates still follow

⁹Achieved in OpenFAST by disabling all rigid-body translational and rotational DOF except for the pitch rotation.

a well-defined, uni-modal distribution, as shown in Figure 5.15. This also means that, unlike in the case of regular waves, the expected value is equivalent to the mean of the ICR coordinates.




Figure 5.14: ICR time history during platform motion in a realistic environmental condition.




Figure 5.15: Histogram of ICR coordinates for a FOWT in a realistic environmental condition.

To study the sensitivity of the ICR to irregular waves parameters, the JONSWAP spectrum has been used to generate sea states of three different peak periods and three different significant wave heights, as per Table 5.2. For each condition, 30 random real-

isations were computed,¹⁰ each including 3600s of signal, sampled at 0.0125 s. $E(x_{ICR})$ and $E(z_{ICR})$ were computed for each random realisation, and then the distributions over the 30 realisations were studied, as shown in Figures 5.16 and 5.17. Similarly to the regular waves test, the focus was on the sensitivity of $E(x_{ICR})$ and $E(z_{ICR})$ to wave frequency (peak period in this case), and to the wave height (or, significant wave height).




Figure 5.16: $E(x_{ICR})$ in irregular waves, for three frequencies and three wave heights. Box plots based on 30 random realisations.




Figure 5.17: $E(z_{ICR})$ in irregular waves, for three frequencies and three wave heights. Box plots based on 30 random realisations.

$E(x_{ICR})$ generally decreases, and $E(z_{ICR})$ increases with frequency increasing, for all wave heights studied. This behaviour is consistent with that observed in the regular waves test, except for the lowest frequency (0.05 rad s^{-1}), where the regular and irregular waves loading resulted in opposite sign $E(x_{ICR})$ and $E(z_{ICR})$.

No statistically significant dependence of ICR on wave height was observed. As can be seen in Figures 5.16 and 5.17, for any frequency, the ranges of the expected

¹⁰The number of random realisations was dictated by the number of samples required to perform the t-test, to study the statistical significance of the trends observed.

value of ICR are overlapping, as also confirmed by the Student's t-test.¹¹ As shown in Section 5.3.1, regular low-frequency waves excite nonlinear hydrodynamic and mooring loads, which highly affect the rotational behaviour of the platform, increasing the sensitivity of the ICR to wave height. However, since an irregular sea state is composed of waves of various frequencies, these nonlinear effects are alleviated by the effect of higher frequency waves, which reduce that sensitivity.

5.4 ICR sensitivity to complex environmental conditions

The study of the ICR of FOWT in waves only provided a limited understanding of the behaviour in real operational scenarios. Therefore, this test aims to establish the dependence of the ICR on more complex environmental loading conditions, to understand the factors driving the ICR, and to establish the applicability of the ICR as a design metric.

5.4.1 Study setup

The sensitivity of the ICR to environmental loading conditions is studied through a series of simulations of gradually increasing complexity, as outlined in Table 5.8.

Table 5.8: Impact of varied environmental loading on the ICR – conditions considered. γ – peak-shape factor.

Group (-)	Wind				Waves				Current
	Type (-)	V_s (m s ⁻¹)	TI (-)	Shear (-)	Type (-)	H_s (m)	T_p (s)	γ (-)	V_c (m s ⁻¹)
A	-	-	-	-	Reg	1.87	5–17.0	-	-
B	-	-	-	-	JONSWAP	1.87	7.47	1.0	-
C	Steady	3–25	-	0.2	Reg	1.87	7.47	1.0	-
D	Steady	11.4	-	0.1–0.4	Reg	1.87	7.47	1.0	-
E	Turbulent	11.4	0.1–0.3	0.2	Reg	1.87	7.47	1.0	-
F	-	-	-	-	Reg	1.87	7.47	1.0	0.5–1.2
G	Turbulent	11.4	0.17	0.2	JONSWAP	1.87	7.47	1.0	0.85

First, a simple test in regular waves is performed for a wide range of wave periods (or frequencies), keeping the wave height fixed (group A in Table 5.8). Subsequently, the

¹¹Two-tail test, $\alpha = 0.05$, $t_{critical} = 2.042$, $t_{value} \in [0, 0.722]$.

response in irregular waves is simulated, keeping the peak period and significant wave height at the previously set values (group B). Irregular sea state includes the component waves at various frequencies, with the highest energy content around the peak frequency and contributions from lower- and higher-frequency waves. The complexity is then increased by including the effect of steady wind (from cut-in to cut-out), which puts the FOWT into a new equilibrium position (group C), as well as the effect of the wind shear, which affects the vertical distribution of the wind load (group D). Group E simulations are run in the rated wind speed condition, this time with turbulent wind inflow of varying turbulence intensities. The turbulent wind is expected to increase the dynamic responses at very low frequencies that are not excited otherwise (the second-order wave loads are not modelled). Group F adds the effect of the subsurface current speed (the power law exponent is kept at a value of $1/7$). Finally, group G represents the conditions likely encountered by an operational FOWT, including the combined effects of the stochastic wind, wave, and current loads.

The results obtained for each group will be presented in turn; all results are given in the global inertial coordinate system.

5.4.2 Results

The simulations conducted under group-A environmental conditions (regular waves) are analysed to understand how the external load period (or frequency) influences the motion behaviour and the ICR of the floating system. As presented in Figure 5.18 and Table 5.9, the mean of the x -coordinate of ICR remains around zero regardless of the wave period, and the mean of the z coordinate approaches zero as the wave period increases. As the wave period increases (which corresponds to a decrease in frequency), the distribution of z_{ICR} becomes wider (higher standard deviation).

Two mechanisms are important: i) the difference in the horizontal velocity's increase at two points on the structure due to the increase in pitch motion as the wave period increases towards 31 s and ii) the higher horizontal velocity at both points due to the increased surge motion as the wave period approaches 125 s.

More specifically, the ICR is computed based on the intersection of the lines perpen-




Figure 5.18: Effect of the regular wave period on the distribution of ICR coordinates (group A).

Table 5.9: The ICR of a FOWT in regular waves (group A).

T_p (s)	Surge		Heave		Pitch		x_{ICR}		z_{ICR}	
	mean (m)	std (m)	mean (m)	std (m)	mean (°)	std (°)	mean (m)	std (m)	mean (m)	std (m)
5.00	-0.09	0.07	0.02	0.00	-0.0723	0.0405	3.55	6.28	-34.52	0.63
7.47	-0.10	0.21	0.02	0.02	-0.0766	0.1169	-1.70	21.49	-31.12	0.61
10.00	-0.09	0.38	0.02	0.05	-0.0619	0.2045	-6.38	39.50	-25.79	1.94
12.50	-0.09	0.51	0.02	0.09	-0.0651	0.2577	-4.77	51.61	-19.08	5.15
15.00	-0.09	0.62	0.03	0.14	-0.0637	0.2967	-6.69	62.63	-13.90	10.08
17.50	-0.09	0.76	0.02	0.16	-0.0649	0.3564	-6.85	70.59	-11.43	17.74

dicular to the velocity vectors at two points on the structure located at $z_A = -120$ m and $z_B = -12$ m. When the wave period approaches the pitch natural period (31 s), the pitch motion increases. Higher pitch motion about a point not coinciding with the midpoint between A and B induces a higher contribution to the horizontal velocity at the point further away from the centre of rotation. With higher pitch motion, the imbalance between the two points becomes more pronounced, and the ICR shifts either higher or lower, depending on the direction of the pitch rotation. This leads to the ICR fluctuating over a wider range.

The second effect becomes more prominent when the wave period is closer to the surge natural period (125 s). In this case, higher surge motion imposes the same increase in horizontal velocity at both points. Figure 5.19 illustrates that, for a given rotational velocity, if a significant horizontal translational velocity is superimposed, the ICR, defined as the point where the normals to the velocity vectors intersect, shifts either higher or lower depending on the direction of the translation and the instant of the

periodic motion (i.e. the phase difference between the translation and rotation).¹² This results in both lower and higher values of the ICR, thereby widening the distributions as was observed.

Additional figures supporting this discussion are provided in Figures I.1– I.2 in Appendix I.




Figure 5.19: The effect of superimposing rotational velocity with (left) horizontal and (right) vertical translational velocities on the position of the ICR.

The mean of the x -coordinate of the ICR remains close to zero across all wave periods, which is consistent with the very small heave motion observed in Table 5.9. As illustrated in Figure 5.19, it is vertical translational velocity that shifts the ICR horizontally away from the central axis; with negligible heave, no such horizontal shift is induced and $\bar{x}_{ICR}$ remains near zero.

Finally, the mean z_{ICR} is consistently negative, indicating that the effective centre of rotation lies below the waterline, and shows a trend toward zero with increasing wave period. As illustrated in Figure 5.19, the direction of the ICR shift depends on the phase relationship between the surge and pitch motions. The fact that $\bar{z}_{ICR}$ is shifted upward suggests that, for the conditions considered, surge and pitch are most often of opposite sign — that is, when the platform pitches in one direction, it tends to translate in the opposite direction — which systematically biases the ICR upward.

The simulations of the group C environmental conditions (steady wind) are analysed to understand how the magnitude of the non-fluctuating load applied at the rotor height

¹²Similarly, larger vertical velocity shifts the ICR horizontally away from the platform.

influences the motion behaviour and the ICR of the floating system.

The results presented in Table 5.10 and Figure 5.20 show that the mean of x_{ICR} increases slightly with wind speed while the mean of z_{ICR} follows a nonmonotonic trend. The dynamic motion amplitudes (standard deviations of surge, heave, and pitch) remain largely unchanged across wind speeds, making it difficult to attribute the trend in $\bar{x}_{ICR}$ to any single physical mechanism with confidence. The nonmonotonic trend in $\bar{z}_{ICR}$ is, however, consistent with the subtle nonmonotonic variation in the heave standard deviation observed in Table 5.10: as illustrated in Figure 5.19, vertical translational velocity shifts the ICR horizontally, and a varying heave dynamic amplitude would therefore induce a corresponding variation in $\bar{z}_{ICR}$. As seen previously for the regular wave cases, the distribution of the vertical coordinate gets wider as the wind speed increases.

Table 5.10: The ICR of a FOWT in steady wind (group C).

V_s (m s^{-1})	Surge		Heave		Pitch		x_{ICR}		z_{ICR}	
	mean (m)	std (m)	mean (m)	std (m)	mean ($^\circ$)	std ($^\circ$)	mean (m)	std (m)	mean (m)	std (m)
0.0	-0.08	0.20	0.02	0.03	-0.06	0.10	-2.59	29.80	-20.25	1.07
3.0	2.46	0.20	-0.01	0.02	0.47	0.11	-1.27	19.78	-31.71	2.65
4.9	6.60	0.20	-0.06	0.02	1.34	0.11	-1.19	19.83	-31.76	6.18
6.7	12.90	0.20	-0.18	0.02	2.60	0.11	-0.29	19.70	-31.11	11.31
8.6	21.34	0.20	-0.42	0.03	4.25	0.11	1.17	19.39	-29.59	18.24
11.4	36.15	0.20	-1.13	0.03	7.54	0.11	4.16	20.05	-25.11	28.31




Figure 5.20: Effect of the steady wind speed on the distribution of ICR coordinates (group C).

The next group of simulations (group D) was conducted with a fixed mean hub-height wind speed of 11.4 m s^{-1} , varying the power exponent of the power law wind profile (shear). While typical offshore values lie around 0.10–0.14 [233], higher expo-

nents can occur under stable atmospheric stratification, within wind farm wakes [234], or at nearshore sites; the full range is explored here to assess sensitivity. As illustrated in Figure 5.21, the distribution of x_{ICR} remains mostly unaffected by the wind profile, as explained by negligible changes to the vertical motion (refer to Table 5.11). However, z_{ICR} exhibits a strong dependence on the vertical distribution of wind. Notably, as the power exponent increases, multiple coupled changes can be observed in the rotor-averaged wind speed, the centre of application of the aerodynamic load (centre of effort - shifts upwards), the mean pitch inclination angle, the thrust coefficient, and the aerodynamic damping.¹³ Depending on the trade-off between these interconnected factors, the shear exponent change leads to changes in the mean and dynamic motions of the system, which subsequently affect the mooring loads and the horizontal velocities of the FOWT, which influence the behaviour of z_{ICR} .

Table 5.11: The ICR of a FOWT in steady wind with shear (group D).

Shear (-)	Surge		Heave		Pitch		x_{ICR}		z_{ICR}	
	Mean (m)	Std (m)	Mean (m)	Std (m)	Mean (°)	Std (°)	Mean (m)	Std (m)	Mean (m)	Std (m)
0.10	2.5	0.2	-0.01	0.02	0.47	0.11	-1.3	19.8	-31.7	2.7
0.15	12.9	0.2	-0.18	0.02	2.60	0.11	-0.3	19.7	-31.1	11.3
0.20	36.2	0.2	-1.13	0.03	7.54	0.11	4.2	20.1	-25.1	28.3
0.25	6.6	0.2	-0.06	0.02	1.34	0.11	-1.2	19.8	-31.8	6.2
0.30	21.3	0.2	-0.42	0.03	4.25	0.11	1.2	19.4	-29.6	18.2
0.35	35.9	0.2	-1.11	0.03	7.57	0.11	4.3	20.0	-25.6	29.0
0.40	36.0	0.2	-1.11	0.03	7.61	0.11	4.1	20.2	-25.5	28.9




Figure 5.21: Effect of the wind shear on the distribution of ICR coordinates (group D).

Figures 5.22–5.23 and Tables 5.12–5.13 demonstrate that the ICR distribution does not significantly change in response to changing turbulence intensity (group E) or sub-

¹³As seen in Figure A.1.

surface current velocity (group F).¹⁴ The distributions at all turbulence levels remain normal, with a slightly increasing standard deviation due to increased horizontal load (higher turbulence intensity leads to a higher aerodynamic fluctuating load for any mean wind speed given). The current load on the slender cylinder of the spar floating structure is insignificant compared to other loads on the system.

Table 5.12: The ICR of a FOWT in turbulent wind (group E).

TI (%)	Surge		Heave		Pitch		x_{ICR}		z_{ICR}	
	Mean (m)	Std (m)	Mean (m)	Std (m)	Mean (°)	Std (°)	Mean (m)	Std (m)	Mean (m)	Std (m)
0	36.2	0.20	-1.13	0.03	7.54	0.11	4.2	20.1	-25.1	28.3
10	36.2	0.25	-1.13	0.04	7.53	0.12	4.1	20.3	-25.2	29.7
15	36.2	0.30	-1.13	0.04	7.52	0.13	3.8	20.3	-24.8	31.3
20	36.2	0.35	-1.14	0.05	7.53	0.14	3.8	20.9	-25.1	33.2
25	36.4	0.39	-1.15	0.05	7.56	0.15	3.9	21.4	-25.4	34.7
30	36.6	0.43	-1.16	0.06	7.60	0.16	3.6	22.1	-26.4	37.6




Figure 5.22: Effect of the turbulence intensity on the distribution of ICR coordinates (group E).

Table 5.13: The ICR of a FOWT simulated in varying current velocities (group F).

V_c (m s^{-1})	Surge		Heave		Pitch		x_{ICR}		z_{ICR}	
	Mean (m)	Std (m)	Mean (m)	Std (m)	Mean (°)	Std (°)	Mean (m)	Std (m)	Mean (m)	Std (m)
0.00	-0.10	0.21	0.02	0.02	-0.08	0.12	-1.6	21.4	-31.1	0.62
0.25	0.39	0.17	0.02	0.02	-0.10	0.10	-1.4	19.4	-31.1	0.60
0.50	2.01	0.20	0.02	0.02	-0.04	0.11	0.0	20.9	-31.2	0.66
0.70	3.94	0.20	0.01	0.02	-0.01	0.12	-0.2	21.4	-31.2	0.82
0.80	5.27	0.20	0.01	0.02	0.01	0.11	-0.3	20.6	-31.2	0.91
1.00	8.41	0.21	-0.02	0.02	0.05	0.12	-0.4	23.2	-31.1	1.23
1.20	12.39	0.21	-0.07	0.02	0.11	0.12	0.0	23.2	-31.2	1.60

¹⁴Notice the small range on the y-axis.




Figure 5.23: Effect of the current speed on the distribution of ICR coordinates (group F).

Finally, Figure 5.24 compares the distributions of the ICR coordinates of a FOWT placed in increasingly complex environmental conditions.




Figure 5.24: Effect of the increased environmental loading complexity on the distribution of ICR coordinates. The regular/irregular wave cases only include regular/irregular wave loads; the steady wind case includes steady wind in addition to irregular waves; the complete environment includes irregular waves, turbulent wind, and current loads.

The mean of z_{ICR} shifts nonmonotonically across conditions, a trend that mirrors the nonmonotonic variation in the surge standard deviation (Table 5.14): as discussed previously, the phase relationship between surge and pitch determines the direction of the ICR shift, and a varying surge dynamic amplitude modulates this effect across conditions. The surge standard deviation is largest in the complete environment due to the additional stochastic loads from turbulent wind. The distribution of x_{ICR} remains centred near zero throughout, consistent with heave remaining small across all conditions.

In the presence of regular waves, the vertical coordinate of ICR is confined to a narrow range. However, as irregular waves, wind, and current are gradually introduced, this range expands considerably. Figure 5.25 provides a more detailed view of the effect of the environment complexity on z_{ICR} range. For each loading condition case, the

confidence interval bounds were calculated using the mean and standard deviation of z_{ICR} for various confidence levels. The ranges displayed in the figure represent the differences between the upper and lower bounds, indicating the intervals within which the true mean z_{ICR} is expected to lie, given the specified confidence level. For any confidence level, more complex environmental loading leads to a broader range. This suggests that the complexity of the loading significantly impacts the variability in z_{ICR} and the potential for extreme values. Importantly, broader ranges are also related to increased uncertainty in predicting the true mean of ICR, which may affect the usefulness of the ICR as a metric in the FOWT design process, as it is less likely to remain near a target location.





Figure 5.25: z_{ICR} range for different confidence intervals.

As reported in Table 5.14, the cases with no stochastic loads (i.e. regular wind and/or steady wind) show a relatively low K-S statistic and a high p-value (these metrics are introduced in Section 5.2.5), meaning that the ICR closely follows a normal distribution in these scenarios. In contrast, the low p-value observed for the stochastic loading cases (irregular waves and/or turbulent wind) suggests that the same conclusion cannot be made regarding the normality of the ICR distribution in those instances. In those cases, the observed distribution of z_{ICR} has significantly increased skewness (more asymmetric) and excess kurtosis (heavier tails and sharper peaks), as illustrated in Figure 5.26.

Table 5.14: The ICR of a FOWT under environmental loading conditions of varying complexity. Each row represents one case of a group listed in Table 5.8.

Condition	Surge		Heave		x_{ICR}				z_{ICR}			
	Mean (m)	Std (m)	Mean (m)	Std (m)	Mean (m)	Std (m)	K-S (-)	p (-)	Mean (m)	Std (m)	K-S (-)	p (-)
Reg. waves (A)	-0.08	0.20	0.02	0.03	-2.59	29.80	0.11	0.58	-20.25	1.07	0.14	0.28
Irreg. waves (B)	-0.08	0.13	0.02	0.02	-1.04	25.52	0.10	0.72	-30.10	7.09	0.16	0.16
Steady wind (C)	36.15	0.20	-1.13	0.03	4.16	20.05	0.07	0.98	-25.11	28.31	0.12	0.46
Complete env. (G)	40.48	0.24	-1.38	0.03	2.28	29.23	0.13	0.34	-26.87	47.36	0.15	0.17

Figure 5.26: Skewness and kurtosis of z_{ICR} distribution for different environmental loading conditions.

5.5 ICR sensitivity to design variables

For any environmental loading condition, the motion response of different FOWT designs differs due to unique hydrostatic, hydrodynamic, mooring, and inertial characteristics.¹⁵ Therefore, this test investigates the design features (or variables) that can be adjusted to achieve the desired ICR “location”.

The study includes four design variables: mooring line length, fairlead vertical position and radius, and platform centre of mass vertical position. These are illustrated in Figure 5.2 and listed in Table 5.15. The anchor vertical position (water depth) is adjusted together with the fairlead vertical position variable to keep a constant anchor–fairlead vertical distance and only vary the point of application of the mooring load on the platform.¹⁶

Note that this set of variables is not exhaustive, and each parameter is varied

¹⁵The impact of design parameters on the system dynamics was studied in [235].

¹⁶Because it is deep water, the relatively small changes in the water depth do not significantly affect the hydrodynamic loads on the platform.

independently. In particular, the fairlead position is adjusted without consideration of the floating platform dimensions, which keeps the effects isolated but neglects the proper physical interface between the fairleads and the platform. The catenary mooring line length is modified without altering the fairlead and anchor positions, resulting in changes to the line angle at the fairlead, to the suspended line profile, and to the portion of the line resting on the seabed.

The broad range of design variables ensures that the study covers the entire design space, including some unconventional but realistic designs showing a wide range of motion responses. Special attention was given to the mooring line length parameter, with the range chosen based on mooring pretension. Specifically, the study examined mooring line lengths that resulted in a $\pm 40\%$ variation in pretension compared to the original design to cover a plausible range of designs.

Note that these design variable perturbations may affect the mean heave (draft) of the floating platform due to changes in the underwater geometry and/or the weight of the structure. This effect could significantly impact the response of the FOWT and may be undesirable, as the draft is a critical design parameter established early in the design process. It is often constrained by factors such as port depth and manufacturing or assembly capabilities. Therefore, an additional step was implemented whenever the underwater geometry was modified: the floating platform was re-ballasted to maintain the original draft. This not only affected the total weight of the system but also influenced the weight distribution and centre of mass.

The design variables were perturbed one at a time, and the change in the ICR statistics (in the global inertial coordinate system) was studied. As introduced in Section 5.2.3, the results for each design variant were averaged over 10 environmental condition bins at the Scottish NE8 site [209], with weights based on bin probability, as outlined in Section 5.2.2. The design variables presented in the figure were normalised to 0–1 for clarity and for unbiased comparison between the variables of different magnitudes. The absolute ranges of the variables are given in Table 5.15 for reference; note that these were purposefully selected to cover a very wide design space.

Based on the design trends presented in Figure 5.27, it is clear that variations in

Table 5.15: Original values and ranges of design variables. All values are in (m).

Design variable	Original	Lower bound	Upper bound
Line length (l)	902.2	891.6	924.1
Fairlead radius (r_F)	5.2	4.0	8.0
Fairlead z (z_F)	-70.0	-120	3.5
Platform CM (z_{CM})	-89.92	-110	-84.29

certain design variables significantly impact the position of the ICR. In particular, the results highlighted the significant influence of the vertical location of the centre of mass on the ICR statistics. This effect was most evident for the z coordinate (lower two plots), where the mean and standard deviation of the ICR showed strong negative and positive trends with the vertical position of the centre of mass, respectively. The vertical position of the fairlead was also found to be an important feature affecting the mean of x_{ICR} .

The vertical position of the centre of mass affects the ICR primarily through its influence on pitch stiffness: a lower centre of mass increases the gravitational restoring moment, reducing pitch motion for a given surge excitation. This effectively increases the ratio of surge to pitch velocity, which — given that surge and pitch are predominantly in phase — shifts the ICR downward according to the superposition illustrated in Figure 5.19, consistent with the observed decrease in $\bar{z}_{ICR}$. The vertical position of the fairlead primarily affects $\bar{x}_{ICR}$, which as established is driven by dynamic heave velocity relative to pitch. Since the anchor position is varied in line with the fairlead, the mooring line angle is preserved, but the point of load application along the spar changes, altering the heave-pitch coupling and consequently the phase relationship between the two motions. As established, it is this phase relationship that determines the direction of the horizontal ICR shift, making the fairlead vertical position an influential parameter for $\bar{x}_{ICR}$. These physical interpretations are proposed as plausible explanations consistent with the observed trends; they are not directly supported by dedicated analyses, and should be treated as hypotheses for future investigation.

The effects of other variables on the ICR metrics were relatively minor. Notably, no clear trends were observed for the standard deviation of x_{ICR} , suggesting that the

variables examined do not significantly influence this metric independently.¹⁷




Figure 5.27: ICR statistics for different one-at-a-time design variable perturbations. Original design values are marked with black outlines. Results are averaged over 10 environmental conditions.

To provide more quantitative insights, a sensitivity analysis was conducted using the data from Figure 5.27 to compute the Spearman correlation coefficient for each pair of design variables and ICR statistics (as detailed in Section 5.2.5). As shown in Figure 5.28, the sensitivity analysis confirmed the earlier observations.

These quantitative results allow for ranking the design variables according to their impact on the ICR statistics, helping to identify the variables that offer the greatest potential for ICR-driven design optimisation. This will be further explored in Section 5.7.

¹⁷The “wiggly” behaviour of the curves in Figure 5.27 arises partly because each data point represents a weighted average over multiple environmental condition simulations. While trends for a single selected condition would appear smoother, the present analysis is carried out in the context of applying the ICR concept to a realistic design scenario. Therefore, the average behaviour across a range of conditions is of greater relevance. All simulations were performed with the same random seed; therefore, the random realisation has no impact on the smoothness of the trends.




Figure 5.28: Spearman correlations between the statistics of ICR and the design variables. Values closer to 1 and -1 indicate a stronger correlation. See Figure J.1 in Appendix J for correlations of all responses.

5.6 Responses sensitivity to ICR

To determine the responses that can be improved by incorporating ICR into the design process, a sensitivity analysis was performed based on the full-factorial design of experiments of the two most impactful design variables identified in Section 5.5 (vertical fairlead position and platform centre of mass), resulting in 64 variations of the floating systems.

The metrics considered included the nacelle fore-aft acceleration (referred to as $NcIMUTAx$ s in tables and figures), blade root out-of-plane and tower base fore-aft bending moments ($RootMyc1$, $TwrBsMyt$), and the most loaded mooring line tension at the fairlead ($FAIRTEN2$), as well as the computed x and z coordinates of ICR (in the global inertial coordinate system). For each pair of responses, distribution statistics and the Spearman rank correlation coefficients were computed, as detailed in Section 5.2.5.

As illustrated in the correlation matrix in Figure 5.29, a significant negative correlation (-0.73 to -0.59) was observed between the standard deviation of the ICR horizontal coordinate and several responses, including nacelle acceleration and both the mean and standard deviation of the tower base and blade root bending moments. A strong negative correlation was also observed between the mooring tension and the mean of x_{ICR} (-0.61 to -0.59). A weaker yet significant correlation (0.43 to 0.46) exists between the mean of x_{ICR} and the mean of the tower base and blade root loads. The standard deviation of z_{ICR} also exhibits a moderate correlation with these loads (0.41 to 0.45). Other statistics, such as skewness and kurtosis, were computed and

analysed; however, they did not yield significant correlations and are therefore omitted for clarity. Similarly, the expected value is not included, as it was found to be very close to the mean of the distributions.




Figure 5.29: Spearman correlations between the statistics of ICR and the loads, obtained by varying the design features. Values closer to 1 and -1 indicate a stronger correlation.

The correlation analysis presented here provides information about the strength and direction of a monotonic relationship. However, it does not give specific information about the magnitude of the change in one parameter in response to a shift in another. Therefore, to assess the potential impact of the ICR on the responses of interest, the correlation results should be interpreted together with the information about the ranges of the responses' statistics, as provided in Table 5.16. By analysing the minimum, median, and maximum values, it can be concluded that the ranges of the responses highly correlated with ICR are sufficiently broad to be regarded as useful design metrics, offering adequate design flexibility.

5.7 Design application prime

The insights gained in this Chapter can be applied to a design scenario. The design case is defined based on the identification of the responses that are most sensitive to the ICR (i.e., those that can be improved by adjusting the ICR) (Section 5.6) and the design variables that have a significant impact on ICR and enable these improvements

Table 5.16: Summary statistics of responses based on the data used for the correlation analysis.

	Unit	Min	Median	Max
NcIMUTAx _s std	m s^{-2}	0.06	0.08	0.12
FAIRTEN2 mean	kN	637.04	1276.43	2977.12
FAIRTEN2 std	kN	0.59	7.80	29.98
TwrBsMyt mean	MN m	-5.58	24.14	100.74
TwrBsMyt std	MN m	3.33	4.35	6.48
RootMyc1 mean	MN m	0.80	3.31	9.50
RootMyc1 std	MN m	0.08	0.19	0.76
xICR mean	m	-8.43	-1.88	4.72
xICR std	m	20.40	32.44	110.56
zICR mean	m	-35.47	-26.60	-18.86
zICR std	m	5.97	13.84	46.28

(Section 5.5). The relevant correlations are summarised in Figure 5.30.




Figure 5.30: ICR case study: design variables considered and their correlation with the ICR and responses statistics. Green/red arrows indicate a positive/negative correlation.

One major area would be particularly worth investigating in the context of design improvements through appropriate ICR “location”. Namely, by maximising the standard deviation of the horizontal coordinate of ICR (i.e. by shifting it towards the positive x direction), one could simultaneously reduce the nacelle acceleration and the tower base and blade root bending moments.¹⁸ Lower nacelle acceleration is desired from the point of view of the wind turbine’s efficient and safe operation; reducing the standard deviation of the tower and blade loads is important for system-wide fatigue damage reduction. However, for these improvements to be possible, it would be necessary to find a set of design variables correlated with the standard deviation of the horizontal coordinate of the ICR and the mean of the ICR coordinates. As can be seen

¹⁸Note that the nacelle acceleration is one of the factors contributing to the tower base bending moment; therefore, these responses are not independent.

in the correlation matrix in Figure 5.28 in Section 5.5, none of the variables considered gives any control over the standard deviation of x_{ICR} . Therefore, this metric cannot be modified in a design scenario.

Instead, the mean of x_{ICR} was seen to be strongly negatively correlated with the mean and standard deviation of the mooring loads and, at the same time, negatively correlated with the fairlead position. Therefore, it seems desirable to lower the fairlead position to increase the mean of x_{ICR} to, in turn, decrease the mean and dynamic tension. However, the mean of x_{ICR} is also positively correlated with the mean of the tower and blade loads; therefore, shifting the ICR towards a position that improves the tension characteristics might increase the tower and the blade loads. This effect can be compensated for by lowering the platform centre of mass to effectively reduce the tower and blade loads.

Therefore, the case study presented in Chapter 6.5 will attempt to improve the mooring loads by varying two design variables: the vertical position of the fairlead and the platform centre of mass, keeping the tower base and blade roots mostly unaffected.

5.8 Conclusion

This section provides a summary of the chapter and outlines the specific contributions to knowledge made.

5.8.1 Summary

The instantaneous centre of rotation (ICR) of floating offshore wind turbines (FOWT) was investigated through two studies:

- A fundamental study: derivation of the methods for ICR identification and analysis of the ICR sensitivity to regular and irregular wave load characteristics,
- An applied study: sensitivity of the ICR to more complex environmental conditions characteristics and FOWT design features.

Fundamental investigation The fundamental study aimed to improve the understanding of the dynamic behaviour of FOWT through an investigation of the ICR. A method to calculate the ICR in the time domain and to analyse its statistics have been proposed. Then, a series of time-domain analyses were carried out in regular and irregular waves to study the behaviour of the ICR, and to establish its sensitivity to the wave frequency and height.

Results showed that FOWT rotates about a point that is not fixed in time or space, and the *instantaneous centre of rotation* term should be used; as a matter of fact, the often used terms *Centre of Rotation/Pitch Centre* imply the existence of a single point fixed in time and space, and, as such, should be avoided.

Both coordinates of the ICR span infinite ranges, however, they tend to concentrate around finite locations. From a statistical point of view, the mean and mode of the ICR distributions are generally different from zero, indicating motion asymmetry: the side of the floater facing the incoming waves experiences a higher amplitude of motion. While the x coordinate of the ICR is distributed symmetrically around its mean and follows a unimodal distribution (single peak), the z coordinate presents multimodal distribution with significant skew. Due to the above considerations, it was proposed to use the most likely point of rotation in given environmental conditions, calculated as the expected value of the ICR.

The expected value of the ICR depends on the wave frequency but is insensitive to wave amplitude, except for the lowest frequency case, primarily due to the mooring loads, both linear and nonlinear in platform offset, and viscous loads, nonlinear in wave amplitude. The mechanism leading to this relationship has been identified: large offsets of the platform in the surge direction in long (low-frequency) waves lead to increased mooring loads, which, towards the end of each half of the surge cycle, affect the heave and pitch displacements and velocities introducing higher harmonics to these responses. Through this surge–heave–pitch coupling, the vertical velocities become highly nonlinear in wave height, hence resulting in the highly nonlinear ICR. This distortion of the heave and pitch motion patterns is attributed to two effects: i) hard clipping (flattening out the peaks of the signal), which introduces odd harmonics, ii)

saturation, or the so-called soft clipping (change of the curvature of the peaks of the signal), which introduces both odd and even harmonics. These two effects are caused by the linear and nonlinear parts of the mooring load, respectively. Similar phenomena are commonly recognised in nonlinear optical/audio signal processing engineering. In the high-frequency limit, the expected value of both coordinates of the ICR tends to be zero. In the low-frequency limit, x_{ICR} and z_{ICR} tend to infinity, which corresponds to a pure translation.

Finally, the ICR excited by irregular waves is sensitive to peak period (or frequency); the trend of $ICR(\omega)$ is the opposite of that observed for the regular waves, exhibiting the opposite sign at low frequencies. The high-frequency limit ICR in irregular waves is identical to the regular waves' case. In an irregular sea state, statistically significant sensitivity to wave amplitude was not observed due to the contributions from multiple component waves of different frequencies and the relatively lower (“averaged-out”) influence of the nonlinear effects.

Applied investigation This part of the study investigated the ICR of FOWT through time-domain analysis. It examined the dependence of ICR on complex environmental loading conditions and design features of the support structure, as well as the relationships between the ICR statistics and key FOWT responses.

The statistics of ICR coordinates under increasingly complex external loading conditions were analysed, isolating the impacts of mean wind speed, shear, turbulence intensity, regular wave frequency, irregular waves, and current speed. The results indicated that the ICR coordinates of FOWTs operating in realistic environmental conditions exhibit significant variability, following a heavy-tailed, sharp-peaked distribution with a high likelihood of extreme values. Notably, at a moderate confidence level of 60 %, the range of the vertical ICR coordinate of the spar-type FOWT was approximately 80 m, indicating that precise placement of the ICR is challenging in complex environments.

To inform practical design applications, design variables that could be adjusted to influence ICR statistics are explored, along with how these statistics correlate with

key responses. The design variables most correlated with ICR were identified as the platform's centre of mass and the vertical position of the fairlead. A full-factorial design of experiments revealed significant correlations between ICR and the responses of interest: nacelle acceleration, tower base fore–aft bending moment, blade root out-of-plane moment, and mooring line tension. These were most significantly correlated with the horizontal ICR coordinate.

Sets of highly correlated design variables, ICR statistics, and responses were selected and used to formulate case study – this will be presented in Chapter 6.5 together with other optimisation cases.

5.8.2 Specific contributions

The following contributions to knowledge were made:

- Demonstration that the FOWTs pitch about an axis that is not fixed in time or space, invalidating a common misconception in the field, and development of analytical and time-domain methods to study the instantaneous centre of rotation.
- Determination of the sensitivity of the ICR to characteristics of environmental loading: mean wind speed, turbulence intensity, wind profile, wave frequency, and current speed, and to design features of the support structure.
- Determination of how to apply new insights to modify the designs of floating support structures to simultaneously reduce loads on the mooring lines, tower base, and blade roots, with the caveat that the ICR, alongside other effects, primarily influences the trade-off between dynamic loads at different locations on the structure.

DOUBLE CHECK FOR "semisubmersible"!

Chapter 6

Hybrid-domain, surrogate-augmented model

This chapter is based on the published paper:

P4 Patryniak K, Collu M, Yeter B, Minisi S, Oneto L, Coraddu A, Hybrid frequency-time domain surrogate-augmented model of dynamics of floating offshore wind turbines. (TBC). *Under review in Ocean Engineering*

Co-authors of this paper contributed to the chapter by training the surrogate models and describing the training process. The sections not independently written by K. Patryniak will be clearly marked throughout this chapter.

6.1 Introduction and previous work

Chapter 3 introduced the baseline coupled MDAO framework for integrated FOWT optimisation, while Chapter 4 established the optimisation problem formulation and associated constraints. While the baseline framework presented in Chapter 3 relied on relatively efficient frequency-domain disciplinary models, the evaluation of the optimisation criteria introduced in Chapter 4 requires higher-fidelity approaches, particularly

for accurate fatigue assessment under long-term environmental loading. At the same time, the computational cost must remain tractable for iterative optimisation.

This chapter addresses this challenge through the development of a hybrid-domain, surrogate-augmented dynamic model. The remainder of this section reviews the specific modelling challenges associated with motion, load, and fatigue prediction, organised according to the modelling paradigm, in addition to those already discussed in Chapter 2.

6.1.1 Efficient physics-based models

For clarity, the physics-based models are discussed first in terms of motion and load response modelling, followed by fatigue assessment considerations.

Motion and load response

Simplified dynamic models have long been used to capture the essential coupled dynamics of FOWT systems [236], and linearised formulations in particular offer clear advantages for iterative optimisation. A comprehensive review of such methods is given in [237] and Chapter 2.3. In what follows, a small number of representative models are discussed to frame the context of existing frequency-domain approaches, noting that this discussion is not intended to be exhaustive or systematic.

Early frequency-domain (FD) approaches primarily treated wind- and wave-induced responses separately. For example, Bachynski [238] modelled wave-driven motions by multiplying the hydrodynamic response amplitude operators (RAOs) with the incident wave spectrum, and obtained wind-driven motions by combining wind RAOs with a thrust spectrum precomputed in the time domain. More integrated FD formulations have since been developed. Hegseth et al. [51] assembled a closed-loop aero–hydro–servo–elastic model by expressing structural and control components in state-space form. Hydrodynamic excitation was obtained from MacCamy–Fuchs theory, damping from linearised Morison’s equation, and added mass from analytical 2D formulations. Aerodynamic loads were computed using a linearised quasi-steady blade element momentum method, and mooring forces were represented using the simplified FD model of [161]. While the aerodynamic predictions compared well with a nonlinear

reference, full-system verification was not demonstrated.

Recent (2020-2026) frameworks incorporate more accurate disciplinary models while retaining computational efficiency. RAFT [239], for example, couples potential-flow and strip-theory hydrodynamics with BEM aerodynamics, quasi-static moorings, and linear turbine control within a unified FD solver. Verification showed generally good agreement with OpenFAST for global motions, though noticeable discrepancies remained in predicted tower-base loads. Overall, FD models have progressed from highly simplified spectral formulations to more comprehensive multi-physics frameworks, improving accuracy while maintaining tractability for design-space exploration. Nevertheless, most published models still lack systematic verification or show notable inaccuracies in load prediction.

One specific challenge in FOWT simulation is the accurate modelling of mooring loads. While platform motions can be adequately reproduced with a quasi-static mooring model [135, 240], the quasi-static approach neglects hydrodynamic drag, inertia, and added-mass forces along the line, causing it to significantly underpredict fairlead tensions, including extreme and fatigue loads relative to experimental measurements [224]. Predicting mooring tensions, therefore, requires a dynamic model, especially in the cases of taut and semi-taut mooring systems, where nonlinear restoring becomes pronounced, and linear model errors can reach unacceptable levels, potentially biasing mooring design [241].

Four main families of time-domain formulations can be distinguished: the lumped mass method (LMM), the finite element method (FEM), the finite difference method (FDM), and multi-body dynamics (MBD), although the last is rarely applied in FOWT analysis [242]. In the LMM, such as MoorDyn [224], the line is discretised into point masses connected by massless springs and dampers, with hydrodynamic forces applied via Morison’s equation at each node. FEM-based models (e.g., [243, 244]) additionally model bending and torsional stiffness. Under wave-frequency excitation, all three perform comparably at low and intermediate frequencies [242]; at high frequencies, only FEM remains reliable.

An intermediate approach is the quasi-dynamic model of Trubat et al. [245], which

retains the static catenary shape but corrects the fairlead tension by updating the distributed vertical force along the line to account for inertia and drag — approaching the accuracy of dynamic models at quasi-static computational cost.

Despite their accuracy, fully dynamic time-domain approaches are computationally demanding in the context of design optimisation. Frequency-domain approaches address this: Larsen and Sandvik [161] proposed early simplified models based on the catenary equations with drag and inertia corrections, including a single-DOF representation of mooring line dynamics. Touzon et al. [246] later linearised a full lumped-mass mooring model for frequency-domain solution, demonstrating adequate tension spectra prediction for catenary moorings. More recently, Abdelmoteleb and Bachynski-Polič [247] developed a frequency-domain LMM, implemented in MoorPy, that delivers comparable fatigue and extreme-tension estimates to time-domain equivalents across catenary and semi-taut FOWT mooring designs at a fraction of the computational cost.

A common validation approach for mooring models is to prescribe the fairlead motion directly, either as a time series or a harmonic displacement, and use it to drive the tensions as the model output, decoupled from any platform hydrodynamic model. This has been used to validate lumped-mass and finite element models against wave tank experiments [224, 240, 242], and to characterise the tension response to harmonic excitation analytically in both time and frequency domains [248, 249].

Fatigue assessment

The simplest spectral-based approaches for fatigue analysis are based on the assumption of a narrow-band stress process. The wind and wave-driven loads exhibit distinct statistical properties, as they act in different (although not completely disjoint) frequency regimes. Depending on the location, the first-order wave loads are concentrated in a specific, higher frequency band of the order of 0.1 Hz (or 10 s period), while the difference frequency second order wave loads and turbulent wind loads typically act over a wider range of frequencies at the lower end of the spectrum of the order of 0.001 to 0.01 Hz (or a period of 100 to 1000 s). While each load source can be considered narrow-band, the total load follows a more complex distribution. Consequently,

narrow-band spectral methods are not suitable when both wind and wave loads are applied simultaneously [250], and, as shown in [251], require corrections and calibration factors.

Dual-narrow-band or wide-band frequency-domain methods are more often applied, as they can account for the complex nature of the combined wind-wave loading. Several studies highlighted the importance of incorporating wind and wave-induced loading in fatigue and reliability analyses of FOWTs. The work of Sorensen [252], Burton et al. [164], and Song et al. [253] showed that the variable aerodynamic loads can significantly contribute to fatigue damage of FOWT. Xu et al. [254] and Wang et al. [255] showed that while wind-induced loading is often the primary contributor to fatigue during operational conditions, high sea states with large wave heights can also lead to substantial fatigue damage. In agreement with this, Zhu et al. [256] showed that while the relative contributions of the wind and waves to fatigue damage depend on the control strategy, wind-induced fatigue tends to dominate in lower wind speed operational conditions, and wave-induced fatigue becomes relatively more important in harsh sea states in the high wind speed (reduced thrust) regime.

Simply summing the fatigue damage from low- and high-frequency responses can lead to underestimation when both sources contribute significantly [26, 105]. Ma et al. [105] propose two practical solutions: the combined-spectrum method, which computes total damage from the root-sum-of-squares stress standard deviation; and the dual-narrow-band method, which applies a correction factor reflecting the relative contributions of the two frequency bands. Appendix D of DNV-RP-C203 [26] also provides a simplified combination rule based on mean zero-upcrossing periods. Beyond these, wide-band spectral methods such as the Dirlik [257] and Tovo–Benasciutti [258] formulations are widely regarded as the state-of-the-art in frequency-domain fatigue estimation. A comprehensive overview was given in [259], which highlighted that these methods were initially derived using Monte Carlo sampling of stress processes, followed by inverse FFT reconstruction of synthetic time series from which empirical cycle distributions were obtained. Their parameters are therefore tuned to reproduce rainflow statistics, but they remain approximations of the underlying time-domain process. Ac-

cordingly, DNV-E301 recommends using the dual narrow-band method and checking a few cases with a more accurate rainflow counting method.

Although computationally efficient, particularly valuable in high-resolution multi-body or FEA structural assessments, the spectral methods are generally conservative [105, 259] and lead to over-engineered structures. In contrast, the rainflow counting method operates directly on the stress time series. It does not require assumptions about the underlying spectral shape, making it the most accurate engineering approach for fatigue assessment [105]. When dynamic responses can be efficiently obtained for selected structural details/locations under the relevant metocean and operational conditions, time-domain rainflow counting offers a balance between accuracy and computational cost, making it highly suitable for fatigue evaluation during early-stage design.

6.1.2 Data-driven models

(This section has been written with significant input of A. Coraddu and L. Oneto - co-authors of P4)

From a methodological standpoint, surrogate modelling encompasses a broad family of regression techniques. Classical approaches include low- and high-order polynomial response surfaces, Gaussian process/Kriging models, and radial-basis-function networks, which provide smooth global approximations and, in the case of Gaussian processes, principled uncertainty quantification [171, 260]. More recently, machine-learning models such as support vector regression, random forests, gradient-boosting ensembles, and deep neural networks have gained prominence due to their ability to handle high-dimensional, nonlinear mappings and very large datasets [261, 262]. Each of these techniques offers different trade-offs between expressiveness, data efficiency, robustness to noise, and interpretability, and their selection is typically guided by problem dimensionality, available data, and computational budget [262].

A typical example directly relevant to FOWTs is the construction of a surrogate model that emulates a high-fidelity aero-hydro-servo-elastic simulator. In this setting, the inputs comprise both design variables (e.g., platform geometry, mooring-system

properties, and controller parameters) and environmental descriptors (e.g., wave spectrum, wind speed, and directionality), while the outputs are key performance indicators such as platform motions, tower-base bending moments, and mooring-line tensions. The training dataset is generated by sampling the joint design–environmental space and running the high-fidelity simulator for each sample, typically at a computational cost on the order of an hour per case. This database is then used to train a data-driven surrogate model that learns the mapping from inputs to system responses. Once trained and carefully validated, the surrogate can be evaluated within its domain of applicability in milliseconds, acting as an efficient emulator of the full nonlinear time-domain model. This enables extensive design-space exploration, sensitivity analysis, and scenario screening that would otherwise be computationally prohibitive, particularly in early-stage FOWT design where thousands of evaluations are required before committing resources to a limited number of detailed simulations.

As reviewed in Section 2.3.7 of the literature review chapter, in the last four years (2022–2025) several studies have explored the application of machine learning techniques—such as response surface methods [145,176], artificial neural networks [177,179], Kriging [178], and XGBoost [180]—to predict FOWT motions, natural frequencies, and mooring loads. To reiterate two important observations of that Section:

- Most existing studies employ surrogate models within stochastic optimisation frameworks to approximate design objectives and constraints, rather than replacing specific computational bottlenecks within the coupled dynamic solution process (with one notable exception [180]).
- The trained models are typically limited to narrow design and environmental envelopes: studies focusing on mooring loads are tied to particular platform designs, while platform motion studies assume fixed mooring configurations, thereby restricting broader applicability.

6.1.3 Chapter outline

Building on the simple coupled model introduced in Chapter 3 and the design problem defined in Chapter 4, this chapter presents an efficient approach to the design evaluation of FOWTs for use in optimisation studies. It provides a method for assessing the metrics across the ultimate, fatigue, and serviceability limit states, as required by a holistic design approach at the conceptual design stage. Progressively more complex dynamic models are introduced, starting with a baseline frequency-domain linear model (Section 6.3.1), followed by a hybrid frequency-time domain model (Section 6.3.2), and, finally, a hybrid-domain hybrid-paradigm model incorporating surrogates for potential coefficients and mooring loads (Section 6.4). Each model is evaluated for accuracy and efficiency, with identified limitations addressed in subsequent models. A demonstration of design optimisation performed with each model is presented in Section 6.5, with the chapter concluding in Section 6.6.

6.2 Methodology

The three increasingly complex coupled models presented in this chapter differ in their modelling paradigms: physics-based (PM) or data-driven (DDM), and modelling domains: frequency domain (FD) or time domain (TD), as summarised in Figure 6.1. Each successive model builds upon the previous one, aiming to achieve an optimal balance between accuracy and computational efficiency. The baseline model (C1) relies entirely on the motion and load response solution in the frequency domain. As will be demonstrated in Section 6.3.1, such a model has limitations in both accuracy, particularly regarding the prediction of mooring tension and tower base bending moment, and efficiency, especially in solving the hydrodynamic problem. The second model introduces two modifications to address these issues. Firstly, the motion response, which is efficiently computed in the frequency domain, is transformed to the time domain (Section 6.3.2). This transformation enables the direct calculation of structural loads in the time domain, allowing for accounting for some motion-load nonlinearities (a method inspired by the work of Pegalajar-Jurado et al. [263]). Secondly, in Section 6.4, two

surrogate models are implemented: one to replace the boundary element method for hydrodynamic coefficient computation, and another for the computation of the mooring loads.

Model	C1		C2		C3	
	Paradigm	Domain	Paradigm	Domain	Paradigm	Domain
Mean motion					PM	
Potential coefficients				FD	DDM	FD
Dynamic motions	PM	FD	PM	FD → TD	PM	FD → TD
Tower loads				TD		TD
Mooring loads					DDM	
Section	3.1		3.2		4.1, 4.2	

Figure 6.1: Dynamic models overview. PM–physical model; DDM–data-driven model; FD–frequency domain; TD–time domain.

6.2.1 FOWT system

The models developed are tested on the example of the OC3 Hywind spar floating platform [149] with the 5 MW reference rotor [264], as detailed in Tables 5.4–5.5 in Chapter 5. For the study of this chapter, however, the mooring system was sized to ensure the mean offset below 10% of the water depth and no vertical load on the anchors in the rated wind (maximum thrust) condition.¹ The characteristics of the modified mooring system are given in Table 6.1.

Table 6.1: Mooring characteristics of the reference FOWT design.

Property	Unit	Value
Fairlead height	m	−70
Fairlead radius	m	5.2
Water depth (anchor position)	m	−320
Anchor radius	m	845.2
Unstretched line length	m	902.2
Chain diameter (nominal)	m	0.07
Pretension	kN	859.0

¹This was achieved by modifying the original line diameter and anchor radius, as one of the possible ways of achieving the desired response. Note that this is not an optimal design in any sense.




Figure 6.2: Reference spar FOWT system illustration.

6.2.2 Verification study setup

Each model is verified against OpenFAST, as detailed in Section 5.2.2. For this verification, OpenFAST was configured to match the capabilities and assumptions of the present framework as closely as possible. Mooring loads were computed using the dynamic lumped-mass model MoorDyn [224]. No current loading was considered, and only co-directional wind and wave conditions were simulated. The baseline variable-speed variable-pitch controller of the NREL 5 MW turbine was enabled, where appropriate (as explicitly stated in the relevant sections) - it allowed for blade pitch and rotor speed variations depending on the mean wind speed and slight rotor speed variations due to the system motion and turbulent inflow. The primary remaining discrepancy between the present framework and OpenFAST concerns the treatment of viscous excitation: the present approach accounts only for viscous damping associated with platform motion, whereas OpenFAST additionally includes viscous excitation induced by water-particle velocities.

To verify the FOWT motions and loads, three series of verification cases are performed: with only the wave load (cases 1. x), with only the wind load (cases 2. x), and with both wind and wave loads (cases 3. x), to provide insight into the mechanisms driving the discrepancies. The analyses rely on the metocean data from the Scottish sectoral marine plan site “NE8” [209]. As summarised in Table 6.2, the wave-only

test cases are run at varying peak periods and significant wave heights, to investigate the response in different frequency and amplitude regimes. Wind-only and wind-wave cases are run in below-rated, rated, and above-rated conditions. The significant wave height H_s and peak period T_p were selected based on the joint wind–wave statistics, as detailed in Sections 4.4.1– 4.4.3.

Table 6.2: Dynamic model verification cases.

Case	Wave-only cases					Wind-only cases				Wind and wave cases		
	1.1	1.2	1.3	1.4	1.5	2.1	2.2	2.3	2.4	3.1	3.2	3.3
H_s (m)	1.0	2.0	1.0	2.0	1.0	-	-	-	-	1.44	1.87	2.68
T_p (s)	8.0	8.0	16.0	16.0	32.0	-	-	-	-	8.14	7.48	7.63
U_∞ (m s ⁻¹)	-	-	-	-	-	4.0	8.5	11.4	15.1	8.5	11.4	15.1

The fatigue assessment in the verification studies is carried out for a single sea state using a 1-hour simulated response, which is then extrapolated to the full lifetime of the structure. This approach contrasts with a fully probabilistic fatigue assessment, where contributions from all sea states expected over the structure’s lifetime are accounted for. By focusing on a single sea state, direct comparisons of fatigue damage across models can be made without averaging over multiple environmental loading realisations. When required, multiple representative sea states are analysed—for example, to investigate fatigue behaviour under different loading regimes.

A single wind and waves direction of 0° is modelled, and a fatigue reduction factor R_μ is applied to approximately account for the directionality of the waves (distribution of the mean wave direction), following the approach proposed in [215]. The metocean data is used to obtain the probability of each mean wave direction $p(\mu_i)$ (shown in Figure 6.3), and the correction factor is obtained according to Equation 4.17.

The output time signals from OpenFAST (tower base bending moment and mooring tension at the fairlead) are post-processed to compute the stress and fatigue damage, which are not direct outputs of OpenFAST (the method will be outlined in Section 6.3.1).




Figure 6.3: Wave rose for the NE8 site.

6.3 Physics-based model

The subsequent sections introduce and verify the baseline frequency domain model (C1) and the modified hybrid frequency-time domain model (C2). The surrogate-supported model (C3) will be introduced in a separate section (Section 6.4).

6.3.1 Baseline coupled model of dynamics (C1)

The baseline dynamic model is conceptually very similar to the model introduced in Chapter 3. However, several important modifications have been implemented, including the introduction of the aerodynamic forcing term. For clarity, the model is therefore presented anew in this section. Where applicable, details that remain unchanged are referenced to the earlier model rather than repeated, while all relevant assumptions and governing equations are restated to ensure completeness and readability.

The baseline model relies on the prediction of the steady-state position of FOWT as excited by mean aero- and mooring loads, as well as the dynamic, harmonic response due to the fluctuating wave and wind loads, both the deterministic component related to the variation of the wind load due to the rotor rotation, as well as the stochastic component related to the turbulent wind inflow.

Static equilibrium

The floating support structure is modelled as a rigid body with three DOFs (surge, heave, and pitch). Only in-plane loads and motion response are considered in the

current study. The static response (equilibrium position) of the floating system is obtained by iteratively solving for the mean aerodynamic and mooring load on the platform ($\bar{\mathbf{F}}_{aero}$, $\bar{\mathbf{F}}_{moor}$) and the offset and attitude of the system ($\bar{\boldsymbol{\eta}}$), considering the hydrodynamic and gravitational restoring ($\mathbf{C}_h$, $\mathbf{C}_g$), using an efficient Broyden solver, as recommended in [196] and Chapter 3:

$$(\mathbf{C}_h + \mathbf{C}_g)\bar{\boldsymbol{\eta}} = \bar{\mathbf{F}}_{aero}(\bar{\boldsymbol{\eta}}) + \bar{\mathbf{F}}_{moor}(\bar{\boldsymbol{\eta}}) \quad (6.1)$$

The load on the platform from the mooring system is modelled using a quasi-static approach implemented in MoorPy [190] (as introduced in Section 3.2.2). The mean aerodynamic load and its derivatives are solved by a steady blade element momentum model implemented in CCBlade [265], including the Prandtl tip and hub loss corrections and wake rotation (non-zero tangential factors).

Dynamic response

The response of the platform to turbulent wind and wave excitation is computed based on the solution of the 3 DOF forced oscillation problem. The complex amplitude of motion $\hat{\boldsymbol{\eta}}(\omega)$ is obtained by solving the equation of motion in the frequency domain:

$$\hat{\boldsymbol{\eta}}(\omega)[- \omega^2(\mathbf{M} + \mathbf{A}(\omega)) + i\omega\mathbf{B}(\omega) + \mathbf{C}] = \hat{\mathbf{F}}_{hydro}(\omega) + \hat{\mathbf{F}}_{aero}(\omega) \quad (6.2)$$

The nomenclature and computation of the left-hand-side terms are provided in Section 3.2.2. The only change is that, in this model, the derivatives of the aerodynamic load, i.e., aerodynamic damping contributing to the $\mathbf{B}(\omega)$ term, are computed by a steady blade element momentum model (CCBlade [265]).

On the right-hand side of Equation 6.2, $\hat{\mathbf{F}}_{hydro}(\omega)$ and $\hat{\mathbf{F}}_{aero}(\omega)$ are the hydrodynamic and aerodynamic fluctuating loads, respectively. More specifically, $\hat{\mathbf{F}}_{hydro}(\omega)$ is the frequency- and direction-dependent complex wave excitation matrix, obtained as a product of the first order wave excitation coefficient $\mathbf{X}(\omega)$ from the diffraction problem solution, and the wave amplitude in the frequency domain $\hat{\zeta}(\omega)$:

$$\hat{\mathbf{F}}_{hydro}(\omega) = \hat{\zeta}(\omega) \mathbf{X}(\omega) \quad (6.3)$$

The complex wave amplitude is derived with amplitudes from the two-parameter JON-SWAP spectrum $\mathbf{S}_J(\omega)$ and random phases $\phi(\omega)$:

$$\hat{\zeta}(\omega) = \sqrt{2S_J(\omega)d\omega} \cdot \exp(i\phi(\omega)) \quad (6.4)$$

$\hat{\mathbf{F}}_{aero}(\omega)$ in Equation 6.2 is the frequency-dependent complex turbulent wind excitation. It is derived based on the fast Fourier transform (FFT) of the fluctuating aerodynamic load time series precomputed using OpenFAST [266] for a range of mean wind speeds, turbulence intensity factors, and rotor plane inclination angles, with all rigid and flexible DOFs disabled (therefore specific to a given rotor, but agnostic to the support structure design). In the simulations, the controller was disabled, and the rotor speed and blade pitch angle were set to the values from the control table, interpolated based on the mean wind speed. Directly precomputing the aerodynamic load and not only the wind inflow velocities was possible, as this study concerns a fixed rotor design. One of the benefits of this approach is that some of the nonlinearities in the wind load could be accounted for in the otherwise linear approach.

Equation 6.2 is used to compute the total complex motion response amplitude $\hat{\boldsymbol{\eta}}(\omega)$. This amplitude is then used to compute the response spectra and statistics. The total motion response power spectrum is computed according to its definition:

$$\mathbf{S}_{\boldsymbol{\eta}}(\omega)d\omega = \frac{1}{2}|\hat{\boldsymbol{\eta}}(\omega)|^2 \quad (6.5)$$

and the standard deviation of the motions is obtained based on the zeroth moment of the response spectrum:

$$\boldsymbol{\sigma}_{\boldsymbol{\eta}} = \left(\int_0^\infty \mathbf{S}_{\boldsymbol{\eta}}(\omega) d\omega \right)^{\frac{1}{2}} \quad (6.6)$$

Loads assessment

The motion response complex amplitudes obtained from the dynamic model are used to compute the system loads, specifically the tower base bending moment and the most loaded mooring line tension at the fairlead. These two locations are considered critical from the perspective of ultimate and fatigue limit states assessment.

The calculation of the tower base bending moment involves summing up the contributions of the aerodynamic load on the rotor and tower, transformed to the tower base-fixed coordinate system, as well as the contributions from the inertial effects of the acceleration and inclination of the structures above the tower base, which include the tower itself and the rotor-nacelle assembly. The computation of the mean bending moment relies on the input of the mean inflow velocity, aerodynamic load, and pitch displacement. In contrast, the dynamic tower load calculation requires the input of frequency-domain amplitudes of the aerodynamic load, pitch displacement, and acceleration, all computed at the mean position of the system.

The mean moment due to aerodynamic drag on the tower is computed based on the integration of the drag along the tower with the drag coefficient c_d and wind shear exponent α :²

$$dM_d(z) = \frac{1}{2} \rho c_d V_{ref}^2 \left(\frac{z}{z_{ref}} \right)^{2\alpha} dA(z) z \quad (6.7)$$

The moment due to the acceleration of the i -th mass item located at the height z_i above the tower base is formulated as:

$$M_{accel,i} = (m_i a_i) z_i \quad (6.8)$$

The moment due to the inclination of the weight item N_i over the tower base is formulated as:

$$M_{incl,i} = z_i (\sin(\theta + \delta\theta) - \sin(\theta)) N_i \quad (6.9)$$

accounting for the reduction of the effective lever of the weight with the mean pitch

²The results presented in this chapter were obtained assuming $\alpha = 0.2$, as recommended by the IEC 61400-1:2019 Section 6.3.2.2 [189] for the normal wind model. However, lower values, 0.11-0.14, are typically assumed and recommended in more recent editions of standards.

angle.

The tower load power spectrum is defined as:

$$S_{M_y}(\omega)d\omega = \frac{1}{2}|M_y(\omega)|^2 \quad (6.10)$$

and the standard deviation of the tower base load is obtained based on the zeroth moment of the spectrum:

$$\sigma_{M_y} = \left(\int_0^\infty S_{M_y}(\omega) d\omega \right)^{\frac{1}{2}} \quad (6.11)$$

As shown in the force-displacement curve in Figure 6.4, mooring tension is nonlinear with respect to the platform offset. Therefore, estimating mooring tension statistics




Figure 6.4: Force-displacement curves and length of chain on sea bed curves for the upwind and downwind mooring lines.

from the frequency-domain solution is non-trivial. In this baseline model, the mean mooring tension was obtained as the static tension in the mean position of the platform:

$$\bar{T} = T|_{\bar{\eta}} \quad (6.12)$$

Due to the nonlinearity of the tension in platform offset, this is expected to slightly underestimate the mean tension, especially at large mean offsets where the tension is most nonlinear. For the computation of the standard deviation of tension, three

methods were considered, all based on the quasi-static model MoorPy [190] with varying inputs:

- Based on the mean and standard deviation of motion, compute static tension in the mean and mean plus standard deviation position. Take the standard deviation of tension as the difference between the two values:

$$\sigma_T|_{\bar{\eta}} = T|_{\bar{\eta}+\sigma_{\eta}} - T|_{\bar{\eta}} \quad (6.13)$$

- For each frequency component of the dynamic motion amplitude $\hat{\eta}(\omega)$, compute the static tension in the mean plus dynamic motion amplitude, and subtract the static tension in the mean offset position; construct the tension power spectrum (only valid for a given mean offset) and integrate it to obtain the tension standard deviation at the particular mean offset.

$$T(\omega)|_{\bar{\eta}} = T|_{\bar{\eta}+\hat{\eta}(\omega)} - T|_{\bar{\eta}} \quad (6.14)$$

- For a given mean platform position, compute the Jacobian of tension with respect to the 6 DOF motions, and multiply it with the dynamic motion amplitude. Construct the tension power spectrum (only valid for small oscillations about the given mean offset), and integrate it to obtain the tension standard deviation.

$$T(\omega)|_{\bar{\eta}} = \left. \frac{dT}{d\eta} \right|_{\bar{\eta}} \hat{\eta}(\omega) \quad (6.15)$$

The first method was highly efficient, requiring only the computation of static tension at two offset points. However, it was found to significantly underestimate the standard deviation of the tension. The other two methods yielded consistent results. Among them, the method based on the tension Jacobian was computationally more efficient: it required only seven evaluations to construct the Jacobian using forward differences, whereas the other method required one evaluation for each frequency component of the motion amplitude, totalling 75 evaluations. Consequently, the current

coupled model (C1) employs the Jacobian-based method (Equation 6.15).

Given the frequency domain tension amplitudes, the tension power spectrum is obtained as:

$$S_T(\omega)d\omega = \frac{1}{2}|T(\omega)|^2 \quad (6.16)$$

and the standard deviation of the tension is obtained based on the zeroth moment of the spectrum:

$$\sigma_T = \left(\int_0^\infty S_T(\omega) d\omega \right)^{\frac{1}{2}} \quad (6.17)$$

Fatigue assessment

The fatigue assessment follows the DNV suggested S-N approach [267], which involves finding the number of cycles n that occurred at critical areas of the structure at a given hotspot stress range and comparing them with the number of cycles to fail N . The number of stress cycles to failure N for each stress range $\Delta\sigma$ is obtained from previous fatigue test results reported in DNVGL-RP-C203 [268] as curves characterised by the intercept of $\log(N)$ axis ($\log(a)$) and the negative inverse slope (m):

$$\log(N) = \log(a) - m \log(\Delta\sigma) \quad (6.18)$$

The fatigue contribution δ_i due to environmental condition i is computed as:³

$$\delta_i = \frac{n_i}{a} E[\sigma_{S,i}^m] R_\mu \quad (6.19)$$

where R_μ is the wave directionality correction factor. The number of stress cycles (n_i) is calculated by multiplying the probability of the i -th environmental state p_i and the design life by the zero up-crossing rate ν_0^+ :

$$n_i = p_i T_D \nu_0^+ \quad (6.20)$$

³Although fatigue damage contribution is typically denoted as d_i , here it is swapped for δ_i as d_i is already used in another context in Chapter 4.

The up-crossing rate is computed from the moments of the stress power spectrum:

$$\nu_0^+ = \sqrt{M_2/M_0} \quad (6.21)$$

The $E[\sigma_{S,i}^m]$ term is the expected value of the stress amplitude raised to the power of m . Although the stress process due to the combined wind and wave loading is expected not to be narrow-band, the narrow-band assumption is made in this simple frequency-domain baseline model. The assumption of the stress being Rayleigh distributed allows for the estimation of the expected value of the nominal stress ranges as:

$$E[\sigma_{S,i}^m] = (2\sqrt{2}\sigma_{S,i})^m \Gamma(\frac{m}{2} + 1) \quad (6.22)$$

The implications of this assumption are discussed in subsequent sections, where later models demonstrate how it can be relaxed. The parameters of the adopted S–N curves are listed in Table 6.3. The C1 curve from DNV-RP-C203 [268] is applied for the tower, assuming good-quality welded joints in air, while the curve for the mooring chain is taken from DNV-OS-E301 [213] for a studless chain.

Table 6.3: Relevant S–N curves parameters.

Location	Material	m	a
Tower base	Welded joint	3.0	2.8119×10^{12}
Fairlead	Studless chain	3.0	6.0×10^{10}

The stress range calculations for the mooring line and tower base section are detailed below.

Mooring lines This analysis focuses on tension-tension fatigue, excluding the out-of-plane bending effects. The nominal stress range (double amplitude) in the i -th environmental condition $\sigma_{S,i}$ is given as the tension range ($\sigma_{T,i}$) divided by the nominal cross-section area of the chain (corrected by corrosion):

$$d_{corr} = d - 0.1 \frac{mm}{year} \cdot 25 \text{ year} \quad (6.23)$$

$$\sigma_{S,i} = \sigma_{T,i}/(2\pi d_{corr}^2/4) \quad (6.24)$$

Tower base The nominal stress at the tower base is calculated based on DNV-GL-RP-C202 recommendation as the difference between the axial compressive and tensile stress:

$$\sigma_{S_i} = \sigma_{S,i,bending} - \sigma_{S,i,comp} \quad (6.25)$$

$$\sigma_{S,i,bending} = \frac{r_i M_{y,i}}{I_{y,i}} \quad (6.26)$$

$$\sigma_{S,comp} = Ma_z \cos(\bar{\theta})/A_i \quad (6.27)$$

Then, the hotspot stress range $\Delta\sigma$ is calculated as the product of the nominal stress and the appropriate stress concentration factor (SCF):

$$\Delta\sigma = \sigma_{S,i} SCF \quad (6.28)$$

The SCF is applied to account for geometric discontinuity (e.g., slope between two sections of varying diameters), complexity of connections, and manufacturing tolerances, in accordance with [26]. The obtained hotspot stress range is then used in the fatigue calculations, which are similar to those outlined for the mooring lines in Equations 6.19–6.22.

Verification

The environmental conditions for verification of the model are outlined in Table 6.2 in Section 6.2.2. To reiterate, cases 1.x relate to wave-only loading, case 2.x relate to wind-only cases, and cases 3.x relate to combined wind-wave loading of increasing severity.

The comparison of the mean and standard deviation of the surge, heave and pitch motions, nacelle acceleration, tension at the fairlead of the most loaded mooring line, and the tower base bending moment, obtained from OpenFAST (OF) and the current baseline (fully frequency-domain) model (C1), is given in Figures 6.5–6.6, with detailed numerical results presented in Tables E.1–E.2 in Appendix E.

As illustrated in Figure 6.5, the mean motion and load responses closely matched the OpenFAST results, with slight overprediction of the offsets (errors below 5.0 %), attributed to the overprediction of the mean aerodynamic load, underprediction of the mooring tension (error below 1.9 %), and overprediction of the tower base bending moment (errors below 7.4 %).




Figure 6.5: Verification of model C1 in operational cases. Mean loads. See Table 6.2 for reference.

As illustrated in Figure 6.6, the prediction of the wave-driven motion response and tower base bending moment (cases 1. x) matched the benchmark results well, with discrepancies generally below 10 %. However, significant discrepancies were seen in the mooring loads. In particular, an underprediction of about 80 % was seen in the standard deviation of the mooring tension induced by the waves of peak period of 8 s (cases 1.1 and 1.2), independent of the wave height. A less severe but still significant overprediction of about 30 % was seen in the cases of peak period of 16 s (cases 1.3 and 1.4), down to 3.4 % in longest 32 s waves (case 1.5). These discrepancies are attributed to two factors: the misprediction of surge motion, and limitations of the applied quasi-static mooring model, including the neglect of mooring-line dynamics and hydrodynamic loading on the lines. As will be shown, the mooring model performed significantly better in the wind-only cases. The nonlinearity between the hydrodynamic loads and platform motion is believed not to play a major role, as the discrepancies in motions remain similar for both significant wave height cases considered.

In wind-only (low-frequency) loading scenarios (cases 2. x), the misprediction of motion and loads remained below 10.0 %, with the highest errors observed near rated wind conditions.⁴ This behaviour is attributed primarily to limitations of the linearised

⁴Although the heave motion amplitude relative error reached 13.1 %, the absolute error was insignif-




Figure 6.6: Verification results in wave-only, wind-only, and operational cases. Standard deviation of motions and loads. See Table 6.2 for reference.

mooring representation at large surge offsets. Unlike wave-only cases, wind loading introduces a significant mean force on the floating system, resulting in a displaced mean equilibrium position. In addition, wind loads contain substantial low-frequency energy that partly overlaps with the surge natural frequency, thereby amplifying the dynamic surge response around the mean position. As illustrated by the force–displacement relationship in Figure 6.7, larger mean offsets are associated with larger variations in mooring tension due to the nonlinear characteristics of the mooring system. Consequently, the validity of representing the mooring restoring behaviour using a constant linearised stiffness diminishes as the operating point moves further into the nonlinear region of the force–displacement curve, leading to increased discrepancies between the benchmark nonlinear model and the present linear formulation.

For the third and final group of cases, while the discrepancies in the mild conditions were low (case 3.1), the maximum thrust and high wind speed conditions led to a significant 9% overestimation of the tower base bending moment (conservative result). The cases of the most significant tower base bending moment misprediction correspond




Figure 6.7: Mooring force-displacement curve obtained from the response to turbulent wind at different mean wind speeds and fixed turbulence intensity of 14 %.

to the largest platform motions (in both pitch displacement and acceleration). In the current approach, the bending moment is calculated using the sine of the pitch motion. As a result, for larger motions, the assumption of linearity, underlying the computation of the standard deviation from the zeroth spectral moment of the bending moment, is violated. Additionally, the overestimation of the standard deviation of the tower base bending moment is partly explained by the overprediction of the aerodynamic fluctuating load.⁵ While both the precomputation for the current model and the OpenFAST solution utilised the same aerodynamic model, the precomputation was performed for a fixed platform; therefore, the coupling between platform motion and aerodynamic loads has not been accounted for. While OpenFAST resolves this coupling at each time step, the current model only partly accounts for it by including the aerodynamic damping term in the frequency-domain equation of motion. Therefore, the discrepancies between OpenFAST and model C1 can be attributed to both the underestimation of the aerodynamic fluctuating load and the underestimation of aerodynamic damping.

The error in the tension standard deviation was highest in the highest wave height case (3.3) at -12.4% . The relatively high error follows from not accounting for the hydrodynamic loads on the mooring lines (these loads are the highest in the most severe sea state considered).

Table 6.4 compares the spectral-based fatigue damage predictions obtained with

⁵Although not presented here, overprediction of the aerodynamic fluctuating load was seen when comparing the OpenFAST and current model's components, similar to the comparison of later Figure 6.11.

the current model (C1) to the rainflow counting-based fatigue predictions derived from OpenFAST load signals for a range of environmental loading conditions (cases 3.1–3.3). For the tower, the bending stress is computed from the tower base bending moment, while for the mooring lines, the stress is computed from the tension at the first node.

Table 6.4: Comparison of the stress and the resulting fatigue accumulation from a single environmental condition extrapolated to a full lifetime.

Case	Tower base				Mooring (fairlead)			
	Stress (MPa)		Fatigue (-)		Stress (MPa)		Fatigue (-)	
	OF	C1	OF	C1	OF	C1	OF	C1
3.1	9.70	13.71	0.59	4.88	5.13	4.94	0.26	5.85
3.2	14.40	20.37	1.68	17.46	9.23	9.46	1.55	44.76
3.3	18.00	25.71	4.45	34.40	8.60	7.42	1.42	21.16

Consistent across both the C1 model and the OpenFAST approach, the fatigue results are significantly higher than the expected value of around 1. This is primarily due to multiple uncertainties in the fatigue assessment process, most notably the choice of S–N curve and the formulation of the stress concentration factor.

There are large discrepancies between the current frequency-domain method (model C1) and OpenFAST results (approximately a factor of 10). This is, to some extent, explained by the misprediction of the stress values, as indicated in Table 6.4 for the tower base. By inspection of Equation 6.22 ($E[\sigma_{S,i}^m] = (2\sqrt{2}\sigma_{S,i})^m \Gamma(\frac{m}{2} + 1)$), any error in stress should lead to an exponentially proportional error in the fatigue estimation, with the exponent of $m = 3$. When comparing the stress and fatigue errors, it becomes apparent that the stress error does not fully account for the fatigue misprediction. Another factor is the invalidity of the narrow-band stress process assumption. The spectral-based fatigue method assumes that the stress is narrow-band and therefore can be assumed to be Rayleigh-distributed, in which case there is a closed-form expression for the stress expected value, and consequently, for the fatigue damage accumulation. To assess the bandwidth of the stress, the narrow-bandness parameter ϵ can be computed, defined as:

$$\epsilon = \frac{\mu_1}{\sqrt{\mu_0\mu_2}} \quad (6.29)$$

where the spectral moments are given by:

$$\mu_n = \int_0^\infty \omega^n S(\omega) d\omega \quad (6.30)$$

For both the tower base and the mooring line (fairlead) stress, the value of ϵ is significantly below 1 (0.49 and 0.6, respectively), indicating that these processes are not narrow-band and violate the assumption, which contributes to inaccurate fatigue prediction by the spectral model.⁶

Overall, the baseline FD model showed deficiencies in predicting the motions and loads in large mean offset and higher sea state conditions. The main factors were the insufficiency of the linear mooring model (stiffness matrix linearised around the mean offset conditions), nonlinearity of the bending moment with respect to the pitch motion, lack of hydrodynamic loads on the mooring lines, too high precomputed aerodynamic load (or, too low aerodynamic damping) that does not account for the platform motion–aerodynamic load coupling in an exact manner. Importantly, due to the violation of the narrow-band assumption, the spectral-based fatigue assessment method yielded results that significantly deviated from the true values, both in terms of absolute values and design trends, rendering it unsuitable for design optimisation with fatigue damage constraints. The following section outlines the development of a time-series-based load model, which aims to address these limitations.

6.3.2 Hybrid-domain model (C2)

The changes made to the baseline model to improve its accuracy, and the new model verification, are detailed in the following two sections.

Dynamic model and fatigue assessment

To improve the capability of the model to represent the loads nonlinear in platform offset, the complex motions in the frequency domain were transformed to the time

⁶The influence of spectral bandwidth on the validity of narrow-band fatigue assumptions and the resulting fatigue prediction errors is quantified and discussed in [269].

domain by performing the inverse Fourier transform (iFFT) of $\hat{\eta}(\omega)$:

$$\eta(t) = \mathcal{F}^{-1}(\hat{\eta}(\omega)) \quad (6.31)$$

The motion time series was used to directly compute the loads in the tower base and mooring lines, as shown in Figure 6.8. The tower-base loads were computed following the same procedure as for the baseline model, but using a time series of the total motion as input rather than the frequency-domain representation. The mooring loads were computed by feeding the motion time series to the quasi-static mooring model (MoorPy). This method computes the mooring loads at each time step, accounting for the mooring tension nonlinearity with respect to the platform motion (although it does not model the two-way motion-load coupling that a purely time-domain model would).




Figure 6.8: The process behind model C2: from frequency-domain motion to time-domain loads.

Using available load-time signals, the rainflow counting method is applied to directly count the stress cycles. A stress range histogram is generated, and fatigue damage accumulation over a single sea state (1-hour simulation) is computed by summing the contributions from k histogram bins, following the Palmgren-Miner rule, augmented by

the load directionality factor R_μ :

$$D_{ss} = \sum_{i=1}^k \frac{n_i}{N_i} R_\mu = \frac{1}{\bar{a}} \sum_{i=1}^k n_i \Delta \sigma_i^m R_\mu \quad (6.32)$$

Finally, the cumulative fatigue damage over the design life of the structure is obtained by summing the products of the one-hour damage, the design life in hours (T_d), and the probability of occurrence of each particular sea state (p_{ss}):

$$D = T_d \sum_{ss=1}^l D_{ss} p_{ss} \quad (6.33)$$

Note that the verification presented in this section only considers one environmental condition at a time ($l = 1, p_{ss} = 1$).

Verification

Model C2 uses the mean and frequency-domain fluctuating motions obtained from the baseline model as its starting point. These motions are transformed into the time domain to compute the tower base and mooring loads. As a result, the two models differ only in their mean and dynamic load predictions, while the motion statistics are identical. The loads results for three operational cases are presented in Figure 6.9 (mean), and Figure 6.10 (standard deviation), and reported in Tables E.3 and E.4 in Appendix E.




Figure 6.9: Verification of model C2 in operational cases. Mean loads.

The mean tension prediction is marginally closer to the benchmark model (OpenFAST) results, due to the hybrid model's (C2) ability to account for the nonlinearity of the loads in the motion response. However, the mean tower base bending moment




Figure 6.10: Verification of model C2 in operational cases. Standard deviation of the loads.

is further from OpenFAST than the fully frequency-domain model (C1).

In the baseline model, two opposing effects influenced the mean prediction: (i) an underestimation arising from computing the mean based on the zeroth moment of the spectrum of a nonlinear process, and (ii) an overestimation due to an overpredicted mean aerodynamic load. These two errors partly cancelled each other out, leading to a slight overestimation of the mean tower base moment. In the hybrid model, the underestimation from the mean calculation is no longer present. As a result, the overestimation from the aerodynamic load dominates, leading to a larger overall overestimation.

Regarding the dynamic response, the mooring tension is overestimated and deviates further from the benchmark compared to the baseline frequency-domain model, resulting in a 13.1 % relative error in the rated wind condition (3.2). However, transforming the motions into the time domain helped improve the accuracy of the tower base bending moment, as it allowed the nonlinear relationship between the load and the pitch motion to be accounted for. For all cases, the error remains below 4 %. As in the baseline model, the discrepancies are attributed to the overprediction of low-frequency, wind-driven loads.

To support the above discussion, Figure 6.11 presents a comparison of the motion and load time series, as well as their corresponding Fourier amplitudes, for the 8.5 m s^{-1} wind case (case 3.1). All motion and load time signals generally agree well with those from OpenFAST, with a slight overestimation of low-frequency surge motion and an underestimation of heave motion in the 0.025–0.05 Hz range, which corresponds to the most energetic part of the turbulent wind spectrum. It is worth noting that the absolute discrepancies in heave motion are very small and have a negligible influence

on the resulting loads.




Figure 6.11: Verification of the responses - time series (left) and Fourier amplitudes (right) for in-plane platform displacements, nacelle horizontal acceleration, line 2 tension at the fairlead, and tower base bending moment.

As expected from the quasi-static mooring model, the mooring tension solution of the C2 model fails to capture the wave frequency load, although it is significantly lower than the low-frequency contribution, which is predicted accurately. The main discrepancy in the tower base bending moment is the mean value, while the dynamic amplitudes match very well as seen in both the time series and frequency amplitudes (also confirmed by comparison to Figures 6.9 and 6.10).

The small errors in the time series of loads had an effect on the fatigue assessment. Figures 6.12–6.14 compare the 1 h mooring tension time series and the rainflow counting results (stress versus number of cycles curves) between model C2 and OpenFAST, for three operational cases (3.1, 3.2, 3.3).

The over- and underprediction of tension in the low and wave frequency range had opposite effects on the fatigue analysis: the signal of model C2 included more cycles of high stress and fewer cycles of low stress. These two effects balanced each other out, resulting in a good general prediction of fatigue (relative errors of -3.4% , 31.3% , and -18.4% in cases 3.1, 3.2, and 3.3). Depending on the level of over- or under-estimation in a particular condition, and the probability distribution of the conditions over the




Figure 6.12: Tension and stress cycles counting for the mooring line (fairlead). Case 3.1 (below rated conditions).




Figure 6.13: Tension and stress cycles counting for the mooring line (fairlead). Case 3.2 (rated conditions).




Figure 6.14: Tension and stress cycles counting for the mooring line (fairlead). Case 3.3 (above rated conditions).

lifetime of the structure, the net effect can either be over- or under-estimation of total fatigue.

Figures 6.15–6.17 compare the 1 h tower base bending stress time series and the rainflow counting result between model C2 and OpenFAST, for three operational cases (3.1, 3.2, 3.3).




Figure 6.15: Bending stress and stress cycles counting for the tower base section. Case 3.1 (below rated conditions).




Figure 6.16: Bending stress and stress cycles counting for the tower base section. Case 3.2 (rated conditions).

Although there are some differences in predicting the tower base bending moment time history (as shown in the left-hand subplots), these are less significant than in the case of the mooring tension, and so model C2 counts the load cycles more accurately (relative errors in fatigue of 9.1 %, 8.7 %, and 6.8 % in cases 3.1, 3.2, and 3.3). Model C2 consistently overestimates the fatigue, which is a desired, conservative behaviour. Unlike for the mooring lines, a higher cycle count is observed at most stress levels (as shown in the right-hand subplots).

Overall, the hybrid frequency/time-domain model (C2) predicts the tower base loads more accurately than the baseline model (C1), without a significant increase in solution




Figure 6.17: Bending stress and stress cycles counting for the tower base section. Case 3.3 (above rated conditions).

time. However, the mooring tension predicted by the model still exhibits considerable error, which propagates into the fatigue analysis and significantly reduces its accuracy. The next section will discuss the development of a data-driven model to address these shortcomings.

6.4 Data-driven model

The limitations of the physics-based model, including excessive computational effort related to the computation of hydrodynamic coefficients and the misprediction of mooring tension, are addressed by introducing two surrogate models into the workflow depicted in Figure 6.18. The development and verification of the surrogate models are detailed in the next two subsections.




Figure 6.18: Physics-based and data-driven models integration. $\mathbf{x}$ - displacement vector.

6.4.1 Potential hydrodynamics surrogate

The computationally expensive boundary element method solver was replaced with a surrogate model that computes the first-order hydrodynamic potential coefficients in the frequency domain based on the input of the spar floating platform geometry features: draft, nine radii along the draft, centre of buoyancy, and system mass.

Dataset generation

This section presents the inputs and outputs of the dataset, as outlined in Figure 6.19, as well as the dynamic model used to map the inputs to the outputs.




Figure 6.19: Database generation process

Dataset input The independent variables, consisting of the floating platform draft and nine radii, were defined over the ranges listed in Table 6.5. For each sampled design, the dependent parameters: mass and centre of buoyancy, were computed. Each design was subsequently evaluated across a non-uniform frequency vector: frequencies from 0.05 rad s^{-1} to 0.18 rad s^{-1} were sampled at an interval of 0.05 rad s^{-1} , frequencies from 0.18 rad s^{-1} to 1.25 rad s^{-1} were sampled at an interval of 0.01 rad s^{-1} , while the range from 0.18 rad s^{-1} to 2.5 rad s^{-1} , where hydrodynamic coefficients gradients are smaller, was sampled at an interval of 0.5 rad s^{-1} . A full factorial design of these parameters would yield $10 \times 7^9 = 403\,536\,070$ potential configurations, each to be evaluated at 113 frequencies. From this large design space, 7 million samples were randomly selected to ensure broad coverage while maintaining computational feasibility.

Table 6.5: Ranges of the inputs of the dataset for potential coefficients surrogate training.

Input	Min	Max	Interval
Independent parameters:			
Draft (m)	75	120	5
9 radii (m)	3	9	1
Frequency (rad s^{-1})	0.05	2.5	Non-uniform (113 items)
Dependent parameters:			
Mass (kg)	$4.90E + 06$	$23.19E + 06$	-
Centre of buoyancy (m)	-80.85	-26.28	-

The platform draft is capped at the reference value of 120 m, as deeper-drafted structures would be challenging to manufacture and deploy given port restrictions. Drafts below approximately 68 m would result in a negative total pitch stiffness C_{55} (the sum of hydrostatic and gravitational contributions), leading to a loss of initial stability (see Figure 4.4). For context, the shallowest spar-type FOWT currently installed, Hywind Scotland, has a draft of about 78 m, while typical designs range between 90 m and 120 m (e.g., OC3-Hywind reference design, Hywind Demo, Hywind Tampen, UMaine spar). Consequently, the draft range is broadly defined as 75 m–120 m. Platform radii are limited to slightly smaller than that of the reference design (OC3 spar) and up to approximately 50 % larger than its maximum diameter, which remains within the range commonly considered feasible for offshore steel structures. As summarised in Table 6.5, this results in a diverse set of platform geometries spanning a wide range of mass and centre of buoyancy values.

Dataset output The output consisted of the diagonal (uncoupled) and the surge-pitch coupled components of the added mass, radiation damping, and wave excitation coefficients, as summarised in Table 6.6.

Dataset characterisation The 13 input features and 14 target coefficients were normalised to zero mean and unit variance prior to training using a standard scaler. Figure 6.20 presents the marginal distributions of the input and output variables, based

Table 6.6: BEM surrogate dataset outputs (hydrodynamic coefficients).

Output	Components	Units	Description
$Ah_{i,j}$	$(i, j) \in \{(1, 1), (1, 5), (3, 3), (5, 5)\}$	kg, kg m ²	Added mass
$Bh_{i,j}$	$(i, j) \in \{(1, 1), (1, 5), (3, 3), (5, 5)\}$	kg s ⁻¹ , kg m ² s ⁻¹	Radiation damping
$\Re(X_i)$	$i \in \{1, 3, 5\}$	N m ⁻¹	Potential excitation (re)
$\Im(X_i)$	$i \in \{1, 3, 5\}$	N m ⁻¹	Potential excitation (im)

on a representative 5% subset of the dataset. Only two radii are shown, as all radius variables exhibit similar distributions. The geometric input parameters (radii and draft) display approximately uniform coverage across their respective ranges. The derived parameters (mass and vertical centre of buoyancy) exhibit approximately normal distributions, whereas the hydrodynamic coefficients display distinctly non-Gaussian distributions due to their nonlinear dependence on platform geometry and wave frequency.

To visualise the sampling uniformity and input–output relationships, pairwise scatter plots of selected inputs and outputs are shown in Figure 6.21. The input variables exhibit broad and uniform coverage across the sampled space, confirming that the random sampling strategy effectively spans the design domain. The outputs display strongly nonlinear relationships with the inputs. The absence of significant clustering or empty regions indicates that the dataset provides adequate coverage for surrogate model training.

Dynamic simulations For each sampled configuration, frequency-domain radiation–diffraction simulations were performed using the BEM solver Capytaine [193], as detailed in Section 6.3.1. The procedure involved meshing according to the defined parameters, computing the dependent variables (mass and centre of buoyancy), running the simulations, and post-processing the results to extract the required hydrodynamic coefficients. Evaluating 7 million samples, where each design was discretised with $\mathcal{O}(10^3)$ panels, required about 800 h on 60 cores of an Intel(R) Xeon(R) Silver 4216 CPU at 2.1 GHz. Since the hydrodynamic linear coefficients are independent of the platform motion, the derived surrogate is broadly applicable to projects with similar platform topologies (independent of environmental conditions, mooring design,




Figure 6.20: BEM surrogate dataset input and output parameters distributions.

etc.). Consequently, the computational cost associated with generating the training dataset represents a one-time offline effort. The surrogate approach therefore becomes increasingly advantageous when reused across multiple optimisation studies or design projects, reducing the effective computational cost per project.




Figure 6.21: BEM surrogate dataset output versus input parameters scatters.

Model development

(This section has been written with significant input of A. Coraddu and L. Oneto - co-authors of P4)

This section describes the data-driven, machine-learning-based model developed to predict the variables reported in Table 6.6 (14 scalar outputs) from the variables listed in Table 6.5 (9 scalar inputs). The model will be constructed using the dataset introduced in the previous section, which provides approximately 7 million examples of the input-output relation.

This problem can be formulated as a multi-output regression task in machine learning, for which many algorithms with associated hyperparameters are available [261,270]. According to the no-free-lunch theorem, no algorithm can be chosen *a priori* without empirical evaluation on the data [271]. Nevertheless, experience and knowledge of the properties of different algorithms can guide some initial *a priori* choices [261,270].

Given the tabular nature of the data under consideration, shallow machine-learning models [270] would typically be preferred. However, algorithms such as kernel methods [272] or ensemble methods [273] cannot be applied here due to their inability to scale to millions of training examples [270]. In this context, neural networks are a suitable choice, as they can handle large datasets by fully exploiting parallel architectures and hardware accelerators, such as state-of-the-art GPUs [274].

For this reason, in this work, we considered a state-of-the-art shallow neural network for tabular data [261]. In particular, we adopted a fully connected architecture composed of: an input layer of dimension 9; a tunable number $\{1, 2, 4\}$ of hidden fully connected layers, each with a tunable width in $\{1, 2, 4, 8, 16\} \times 10^3$ neurons and activation function chosen from $\{\text{ReLU}, \text{sigmoid}\}$; and a final fully connected layer with linear activation producing the 14-dimensional output. The network was trained using the Adam optimiser with a tunable learning rate in $\{0.001, 0.002, 0.005, 0.01\}$ for a tunable number of epochs in $\{5, 6, \dots, 100\}$, minimising the mean squared error between predicted and actual values. It is important to note that, given the extremely large size of the dataset, which required several days of training on an NVIDIA A100 GPU

equipped with 40 GB of RAM, we could not explore a larger set of architectural or training hyperparameters.

In order to tune the free hyperparameters, namely the number of hidden layers, the number of neurons in each hidden layer, the activation functions, and the learning rate, and to estimate the performance of the final model, we employ a state-of-the-art hold-out resampling strategy [262]. Given the very large size of the dataset (approximately 7 million examples), more advanced approaches such as cross-validation are unnecessary and, in any case, computationally impractical. The approach is based on a simple yet robust idea [262]: for each iteration, the original dataset is split without replacement into three independent subsets. These subsets are the training set (approximately 6 million samples), used to learn the model; the validation set (approximately 0.5 million samples), used to tune the hyperparameters; and the test set (approximately 0.5 million samples), used to assess the performance of the final model.

In this work, it is not only important to achieve high accuracy but also to ensure fast prediction times. This implies that smaller networks yield faster inference, although often at the cost of reduced accuracy [261]. For this reason, we search for the smallest (i.e., fastest) network whose mean squared error increases by no more than 5 % with respect to the best-performing network in the validation search space. This 5 % threshold was chosen empirically, as it provides a good trade-off between accuracy and prediction speed.

On the test set, performance can be evaluated using both quantitative metrics [275] (e.g., prediction time, mean squared error, mean absolute error, or coefficient of determination) and qualitative [276] (e.g., scatter plots, histograms of the errors, or trends with respect to individual inputs).

Model verification

To assess the predictive accuracy of the surrogate model across all outputs, a combination of quantitative metrics and graphical diagnostics was employed. First, a range of aggregate error metrics is analysed to provide concise numerical benchmarks of the model’s accuracy across the entire dataset, as detailed in Table 6.7.

Table 6.7: Performance metrics used for surrogate model evaluation. y_i and $\hat{y}_i$ – true and predicted values of the output, $\bar{y}$ – mean of the true values, N – number of samples.

Metric	Description	Formulation
MSE	Average squared deviation between predicted and true values; sensitive to large errors.	$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$
RMSE	Square root of MSE; expressed in the same units as the output variable.	$\text{RMSE} = \sqrt{\text{MSE}}$
MAE	Average absolute deviation; less sensitive to outliers than RMSE.	$\text{MAE} = \frac{1}{N} \sum_{i=1}^N y_i - \hat{y}_i $
R^2	Indicates how much of the variance in the true values is explained by the model; values closer to 1 denote better agreement.	$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}$
NRMSE	RMSE normalised by the range or mean of the true values; provides a scale-independent accuracy measure.	$\text{NRMSE}_{\text{range}} = \frac{\text{RMSE}}{y_{\text{max}} - y_{\text{min}}}$
		$\text{NRMSE}_{\text{mean}} = \frac{\text{RMSE}}{\bar{y}}$

The metrics values are reported in Table 6.8, grouped by the mode (surge, heave, pitch, and surge-pitch coupled mode). The surrogate model achieved near-perfect agreement with the reference hydrodynamic data, with coefficients of determination (R^2) exceeding 0.999 across all motion modes (Table 6.8). The range-normalised errors ($\text{NRMSE}_{\text{range}} < 0.6\%$) and mean-normalised errors ($\text{NRMSE}_{\text{mean}} < 5\%$) indicate excellent consistency throughout the training domain.

Table 6.8: BEM coefficients surrogate aggregate error metrics grouped by mode.

Metric	Surge	Heave	Pitch	Surge-pitch
MSE	3.01e+09	2.71e+07	1.39e+17	3.3e+13
RMSE	3.46e+04	4.88e+03	1.87e+08	4.12e+06
MAE	3.08e+04	3.08e+03	1.71e+08	3.79e+06
R^2	0.999	0.999	0.999	0.999
NRMSE (range)	0.00417	0.00409	0.00456	0.00568
NRMSE (mean)	0.0161	0.0448	0.0188	0.0168

To visualise local variations in model performance, scatter plots of true versus predicted values were generated across different modes: surge, heave, pitch, and surge-pitch coupled mode (Figure 6.22). Perfect agreement would correspond to points lying on

the 45° line. These plots provide an intuitive indication of the global goodness of fit and reveal any systematic trends or bias. To complement this view, histograms of the absolute relative errors are shown alongside the scatter plots. The plots were generated based on 5% randomly selected samples. Frequency was included as one of the model inputs, alongside the platform parameters, meaning each point in the plot corresponds to the prediction for a specific spar design at a specific frequency.

On average, the predictions align closely with the true values, except for the excitation coefficient (X) and radiation damping (B_h) in the heave direction, where significant deviations from the ideal fit are observed. $X_{im,3}$ and $B_{h,33}$ were therefore further investigated through additional diagnostics, including a set of plots that examine the residual behaviour and distributional agreement between the true and predicted values. Specifically, the residuals versus predicted plot highlights potential bias or variance trends in the model errors, the residual histogram shows their overall distribution and symmetry, the kernel density estimates (KDEs) compare the statistical distributions of true and predicted responses, and the quantile–quantile (QQ) plot assesses the normality of residuals.

These diagnostics are presented for B_{33} in Figure 6.23. The residual distribution (top-right panel) shows a small systematic bias, with the mean slightly shifted below zero, indicating a mild tendency of the surrogate to underpredict the radiation damping coefficient $B_{h,33}$. A subtle skewness is also visible, with a heavier left tail and lighter right tail, suggesting that negative residuals (underpredictions) occur more frequently or with larger magnitude. This behaviour is consistent between the residual histogram and the QQ plot, where deviations from the reference line at the lower quantiles further confirm the asymmetry of the residual distribution. Nevertheless, these effects are small in magnitude and do not significantly compromise the overall fidelity of the surrogate.

In terms of $X_{im,3}$, the model exhibits strong predictive performance overall, with residuals narrowly distributed around zero and no evident trend across the predicted range. The true and predicted distributions are nearly identical, confirming that the model accurately captures both the mean behaviour and variability of the data. A slight positive shift in the residual mean and a modest right-tail skew indicate a minor




Figure 6.22: BEM surrogate true-predicted scatter plots with absolute relative error histograms.




Figure 6.23: Diagnostic plots for $B_{h,33}$ prediction.

overprediction tendency, consistent with the residual histogram and QQ plot. Despite these small deviations, the model's errors remain well within acceptable limits.




Figure 6.24: Diagnostic plots for $X_{im,3}$ prediction.

Additionally, the performance of the surrogate model across the full frequency range for each hydrodynamic coefficient was evaluated using the normalised root-mean-square error (NRMSE). This metric was chosen for its dimensionless and scale-independent nature, its robustness in the presence of near-zero or sign-varying values, and its intuitive interpretation, as it expresses the error relative to the mean observed value. To illustrate the model's ability to reproduce design trends, the NRMSE was plotted against one design variable (e.g., draft) while keeping the other (e.g., waterline radius) fixed, as shown in Figures 6.25 and 6.26.




Figure 6.25: NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against waterline radius for different values of draft.




Figure 6.26: NRMSE between the physics-based and data-driven BEM model results (hydrodynamic coefficients), computed over all frequency components. Trends against draft for different values of waterline radius.

An acceptably low NRMSE was observed for most coefficients across the design space, except for the heave components of radiation damping and the real and imaginary parts of wave excitation (B_{33} , $X_{re,3}$, $X_{im,3}$), particularly in designs with a small waterplane area (radius). However, these discrepancies occur in coefficients that contribute minimally to the total response of the FOWT and therefore do not compromise the model’s overall usefulness. Specifically, for the slender spar-type platform, potential damping is small compared to additional linear⁷ and viscous damping, and wave excitation is small compared to aerodynamic excitation.

Nonetheless, to assess the impact of these discrepancies, the motion and load responses of the FOWT were evaluated in three operational conditions, as shown in Figure 6.27 and Table 6.9, by comparing the baseline physics-based frequency-domain model with the version using the BEM surrogate. The resulting errors in total response were generally very small, confirming the validity of the surrogate implementation. Notably, the total solution time was significantly reduced—an essential factor for effective design optimisation.⁸




Figure 6.27: Comparison of the baseline FD model and the FD model with BEM surrogate in terms of the standard deviation of motions and loads. Performed for a reference design redefined with a reduced number of radii. See Table 6.9 for numeric results.

6.4.2 Mooring surrogate

As illustrated in the flow diagram of Figure 6.18, the mooring dynamics are solved at three key stages of the coupled solution process:

⁷ Additional linear damping is added to account for the contributions to hydrodynamic damping other than the linear radiation damping and nonlinear viscous drag, as observed in free-decay tests [149]

⁸ The solution times of all models will be compared in Table 6.5 in Section 6.4.2.

Table 6.9: Relative errors due to incorporating the BEM surrogate model in the frequency domain baseline model (C1).

Case	Platform motion (%)					Loads (%)	
	Surge	Heave	Pitch	Nac. Accel.		T2	My
3.1	1.00	0.52	0.08	-2.25		1.24	-0.38
3.2	1.37	-2.77	-0.15	-4.27		1.71	-0.39
3.3	1.39	-0.99	0.03	-3.23		2.00	-0.38

1. Static response: solution of the mean motions of the platform requires calculating the mooring load on the floating platform. This is achieved using the quasi-static model MoorPy (see Section 6.3.1).
2. Dynamic response: solution of the dynamic motions of the platform around its mean position relies on the mooring stiffness computed at the mean position, which is also done using MoorPy.
3. Loads: given the platform motion, line tensions are determined using either a dynamic model (e.g., MoorDyn) or a quasi-static model (e.g., MoorPy).

In the following, the application of a machine learning method to predict mooring tension at the final stage is presented, addressing the inaccuracies of the physics-based model while avoiding the high computational cost of the time-domain dynamic simulations. The surrogate model maps mooring design parameters and fairlead motion time series to tension time series. The remainder of this section describes the database generation, model training, and verification process.

Dataset generation

A state-of-the-art numerical model is employed to obtain the FOWT motion responses and mooring load time series, thereby generating the synthetic database necessary for surrogate model training in the absence of experimental data. The database generation process consists of the following stages (also illustrated in Figure 6.28):

1. Design of experiments (DOE): define a parameter space covering floating platform and mooring design variables, as well as environmental parameters, using a full

factorial design of N samples.

2. Sampling: randomly select a subset of $M \ll N$ DOE samples,
3. Simulation: perform dynamic simulations for the selected cases,
4. Post-processing: conduct convergence checks and sanity checks on the simulation results,
5. Feature and target selection: identify the parameters and responses to be used as features and targets, drawing from both DOE variables and dynamic simulation outputs.




Figure 6.28: Mooring surrogate database generation process.

Dataset input Every mooring system in this database consists of three catenary lines of equal length, each made entirely of studless chain. The material and hydrodynamic properties are summarised in Table 6.10. Based on these properties and the nominal chain diameter d_{moor} , the linear mass density of the lines m_{moor} and the extensional stiffness EA are computed based on the empirical expressions from [208]:

$$m_{moor} = m_2 d^2 \quad (6.34)$$

$$EA = EA_2 d^2 + EA_3 d^3 \quad (6.35)$$

Keeping the material and hydrodynamic chain properties fixed for all designs, six parameters are needed to uniquely define each variant of the mooring design: nominal

Table 6.10: Mooring chain properties

Property	Value	Unit
Grade	$R3$	-
Mass density per diameter squared (m_2)	2e4	$\text{kg m}^{-1} \text{m}^{-2}$
Stiffness per diameter squared (EA_2)	85.6e9	N m^{-2}
Stiffness per diameter cubed (EA_3)	-3.93e7	N m^{-3}
Normal drag coefficient	1.333	-
Axial drag coefficient	0.639	-
Normal added mass coefficient	1.0	-
Axial added mass coefficient	0.5	-

diameter of the chain d_{moor} (i.e., the diameter of the rod used to produce the chain), line length l_{moor} (measured along the line from fairlead to the anchor), fairlead z -coordinate z_F (point of attachment of a mooring line to the platform), fairlead radius r_F , anchor z -coordinate z_A (equal to the water depth), and anchor radius r_A , all illustrated in Figure 6.29.




Figure 6.29: Mooring line parametrisation - parameters of the design of experiments and ML features.

Design of experiments (DOE) The design space is sampled based on seven scalar variables representing the geometric properties of the platform and the mooring design, as well as four scalar variables representing the environmental loading conditions, as listed in Table 6.11. Full factorial DOE is used to sample the space evenly, resulting in 25 600 design configurations and 108 environmental conditions (2 764 800 combinations

Table 6.11: Primary platform, mooring system, and environmental parameters (design of experiments).

	Min	Max	Units	Count
Platform and mooring design variables				
Nominal chain diameter (d_{moor})	0.056	0.084	m	5
Fairlead radius (r_F)	4.16	6.24	m	4
Fairlead z -coordinate (z_F)	-84.0	-56.0	m	4
Anchor radius (r_A)	676.16	929.72	m	4
Anchor z -coordinate (z_A)	-330.0	-260.0	m	4
Line length (l_{moor})	721.76	1082.64	m	5
Platform vertical centre of mass (z_{CM})	-98.91	-71.93	m	4
Environmental conditions variables				
Mean wind speed (U)	3.0	25.0	m s ⁻¹	4
Turbulence intensity (TI)	12	16	-	3
Significant wave height (H_s)	0.5	6.0	m	3
Peak period (T_p)	6.0	10.0	s	3

in total). Out of these, a random selection of 300 000 samples is performed to limit the computational effort of database generation.⁹ This is an initial set – if the model performance is not satisfactory, additional samples might be added to the training dataset.

Initially, the parameter ranges were intentionally broad to encompass a wide range of possible designs, thereby enhancing the model’s general applicability (see Table F.1 in Appendix F). However, such wide ranges, even with coarse granularity, result in a prohibitively large number of samples to compute. If only a small subset of the broad parameter space can be simulated, the resulting dataset lacks the resolution needed to capture the nonlinear dependence of line tension on the input parameters and is therefore insufficient to train a surrogate model with the required precision. Preliminary verification results in Appendix B illustrate this curse of dimensionality. Hence, the present study employs narrower parameter ranges. While this reduces the general applicability of the surrogate model,¹⁰ it enables the development of a valid

⁹Further to that, some simulations will fail to converge due to an ill-defined system, conflicting variables, or too harsh an environmental loading; therefore, the set of valid samples for training purposes will be smaller than 300 000. In the current case, 170 000 samples were used.

¹⁰The surrogate model is not meant to extrapolate well.

proof-of-concept. The input parameter ranges were therefore based on $\pm 20\%$ of the original design's values, except for the anchor radius, limited to $\pm 10\%$ of the original design's value.

Features extraction The DOE described above serves to generate a diverse dataset of floating systems and their dynamic responses. However, the features used in the machine learning model are not identical to the design variables mentioned earlier. Instead, we select a specific set of features that directly influence the mooring line loads. Given the fixed properties listed in Table 6.10, the catenary mooring design can be uniquely represented by four features: line diameter, line length, radial fairlead-anchor distance, and vertical fairlead-anchor distance, which render the model independent of the platform design. Additionally, two features represent the time series of fairlead motion in the horizontal and vertical directions, making the model independent of environmental conditions. All six features are summarised in Table 6.12.

Table 6.12: Mooring load surrogate model features.

Feature	Min (m)	Max (m)	Type
Nominal chain diameter (d_{moor})	0.056	0.084	Scalar
Line length (l)	721.76	1082.64	Scalar
Fairlead-to-anchor height (z)	204.0	246.0	Scalar
Fairlead-to-anchor radius (R)	672.0	923.48	Scalar
Horizontal motion at fairlead (x_F)	(-)	(-)	Time series 3600 s at 0.05 s
Vertical motion at fairlead (z_F)	(-)	(-)	Time series 3600 s at 0.05 s

Dataset output For the selected samples, numerical simulations are performed to evaluate the dynamic behaviour of the FOWT system under combined wind and wave loading. The analysis focuses on the horizontal and vertical motion of the fairlead, as well as on the corresponding horizontal and vertical loads in the most heavily loaded mooring line, as summarised in Table 6.13. The loads are evaluated at the first chain link connected to the fairlead.

Dataset characterisation Figure 6.30 shows the marginal distributions of the input and output variables, based on a representative 5% subset of the full dataset.

Table 6.13: Mooring load surrogate model targets.

Target	Units	Type
Upwind line F_x at fairlead	kN	Time series 3600 s at 0.05 s
Upwind line F_z at fairlead	kN	Time series 3600 s at 0.05 s

The shown inputs include four geometric parameters (chain diameter, length, and the fairlead-to-anchor vertical and radial distances) together with the mean and standard deviation of the fairlead motions. The shown outputs comprise the mean and standard deviation of the line load components. Overall, the dataset spans a wide range of mooring designs and motion conditions. The multimodal or skewed distributions reflect the discrete sampling of design parameters and the nonlinear nature of the mooring responses. All outputs show broad, asymmetric distributions. The relationships between




Figure 6.30: Mooring surrogate dataset input and output parameters distributions.

the input and output parameters are illustrated in Figure 6.31. For clarity, only one geometric input variable is shown, as the discrete sampling of the design space results in

clustered input–output patterns that are not particularly informative. Importantly, the plots highlight the nonlinear relationships between motion statistics and the resulting tension and fatigue responses.




Figure 6.31: Mooring surrogate dataset output versus input parameters scatters.

Dynamic simulations All simulations are performed using a model of the spar-supported 5MW wind turbine [149, 264], as detailed in Section 6.2.2. However, the features selected for the surrogate training are platform-design independent. Therefore, the choice of the floating platform concept and specific rotor does not affect the applicability of the surrogate to other designs.

The samples are evaluated using OpenFAST. Unlike the verification studies presented earlier in this chapter, the simulations performed for dataset generation employ a less simplified modelling setup. Although flexible structural modes remain disabled, the NREL 5 MW baseline variable-speed, variable-pitch controller is enabled, allow-

ing blade pitch and rotor speed to vary with the mean wind speed, as well as limited rotor-speed fluctuations induced by platform motion and turbulent inflow [264]. The aerodynamic solution is based on blade-element momentum theory combined with the Beddoes–Leishman dynamic stall model, tip- and hub-loss corrections, and tower-shadow effects, as implemented in AeroDyn [277].

Model development

(This section has been written with significant input of A. Coraddu and L. Oneto - co-authors of P4)

This section describes the data-driven, machine-learning-based model developed to predict the variables reported in Table 6.13 (two time series of 72 000 elements) from the variables listed in Table 6.12 (four scalars and two time series of 72 000 elements). The model will be constructed using the dataset introduced in the previous section, which provides approximately 170 000 examples of the input-output relation.

In this section, we do not repeat the concepts presented in Section 6.4.1, as most of them apply here in exactly the same way. The only difference here lies in the machine learning model that we chose to adopt based on the available data. In this case, we have a huge input (144 000 elements) and output (144 000 elements) space, which, moreover, are not simply scalar (four scalar inputs) but sequences (two input sequences and two output sequences), meaning that there is structure both in the input and in the output. Moreover, we have a large number of examples of this input–output relation (170 000 examples), representing almost 100 GB of data.

To handle this amount of data, analogous to what was done in Section 6.4.1, the only viable option is to use a neural network [261]. Because of the nature of the input and output, the most reasonable choice would be to use architectures specifically developed for sequence data, such as Gated Recurrent Units [278], Long Short-Term Memory networks [279], Temporal Convolutional Networks [280], Transformers [281], or Time Series Foundation Models [282]. Nevertheless, while attempting to rely on these types of networks, it soon became apparent that either their accuracy was too low or the size

of the network (with associated memory and computational requirements) was too large to be handled, even with our hardware accelerators, while still achieving a prediction time suitable for subsequent use in optimisation.

As a consequence, we had to adopt a different approach, one that relied on the physical meaning of these inputs and outputs. The input series and a few scalar parameters represent the system and the excitation applied to it, which produce the output sequences. For this reason, given that the structure cannot have infinite memory, the output at time t can only be related to the scalar inputs plus the input series within a window around t , namely $[t - \Delta, t + \Delta]$, with Δ measured in seconds. Moreover, since the structure (FOWT characteristics) does not change over time, the response should not depend on t itself. Thanks to this observation, we only need to construct a neural network applied at time t whose inputs are the four scalars plus the elements of the input series in the window $[t - \Delta, t + \Delta]$, where Δ is tunable in $\{0.1, 0.5, 1, 2, 4, 8\}$ seconds. We then use a neural network exactly like the one adopted in Section 6.4.1, and as output, we predict the values at time t of the two output time series. Despite this reduction in dimensionality, the resulting dataset remained extremely large and required several days of training on an NVIDIA A100 GPU with 40GB of RAM. Consequently, we were unable to explore a broader range of architectural or training hyperparameters, nor larger values of Δ .

In order to tune the free hyperparameters, Δ together with the same hyperparameters of the neural network described in Section 6.4.1, namely the number of hidden layers, the number of neurons in each hidden layer, the activation functions, and the learning rate, and to estimate the performance of the final model, we follow the same procedure described in Section 6.4.1 where 80% of the data were used for training, 10% for validation, and 10% for test.

Model verification

To assess the predictive accuracy of the surrogate model across the two targets, a combination of quantitative metrics and visual diagnostics was employed. The statistics of the force components (mean and standard deviation) are used as scalar values

representing the mooring load time series, where necessary.

First, a table of aggregate error measures, including MSE, RMSE, MAE, MAPE, R^2 , and normalised RMSE,¹¹ is reported to provide concise numerical benchmarks of the model's accuracy across the entire dataset (Table 6.14). The surrogate model

Table 6.14: Mooring model surrogate aggregate error metrics.

Metric	F_x mean	F_x std	F_z mean	F_z std
MSE	6.3774	2.0043	2.5979	3.1995
RMSE	2.5253	1.4157	1.6118	1.7887
MAE	1.3005	0.8445	0.9065	1.0096
MAPE (%)	0.4045	9.9621	0.1699	14.5947
R^2	0.9998	0.9767	0.9999	0.9002
NRMSE (range)	0.0025	0.0251	0.0015	0.0431
NRMSE (mean)	0.0088	0.1157	0.0032	0.2116

demonstrated excellent predictive performance for all response components, with $R^2 > 0.90$ across metrics. Errors in the mean force predictions were negligible (NRMSE < 0.003), indicating near-perfect agreement with the reference model. While the standard deviation predictions showed higher relative errors (up to about 15% in MAPE), the overall accuracy remains high, confirming that the surrogate effectively captures both the mean and variability of the mooring line forces.

To visualise local variations in model performance, scatter plots of true versus predicted values were generated, as introduced in Section 6.4.1, presented in Figure 6.32. The plots confirmed the near-perfect agreement between predicted and true values for the mean components of the mooring loads, with points lying almost exactly along the diagonal line and very low absolute relative errors. In contrast, the standard deviation predictions exhibit a broader scatter, particularly for the vertical load component, indicating comparatively lower accuracy in capturing the variability of the response. Nevertheless, the overall distributions of absolute relative errors remain concentrated near zero, demonstrating that large deviations are rare and that the surrogate model performs consistently well across all components.

¹¹MSE: mean squared error; RMSE: root mean squared error; MAE: mean absolute error; MAPE: mean absolute percentage error; R^2 : coefficient of determination




Figure 6.32: Mooring surrogate true-predicted scatter plots with absolute relative error histograms.

To further examine the structure and distribution of model errors, Figure 6.33 presents the residual-based diagnostic plots for the least accurate prediction case, namely the standard deviation of the vertical mooring load component, as introduced in Section 6.4.1.




Figure 6.33: Diagnostic plots for standard deviation of F_z .

The residuals exhibit a funnel-shaped pattern with increasing variance at higher predicted values, indicating mild heteroscedasticity and slightly reduced accuracy for larger responses. The residual distribution is approximately centred around zero but shows heavier tails, suggesting a few larger deviations in both directions. The kernel density comparison shows that the surrogate accurately reproduces the bimodal nature of the true response distribution, albeit with minor discrepancies around both peaks. The QQ plot reveals noticeable departures from normality in the tails, consistent with the presence of outliers, but the majority of the residuals align well with the theoretical quantiles. Overall, the model accurately captures the general response behaviour, with only limited bias and acceptable error dispersion across the prediction range.

The preceding analysis focused on the statistical characteristics of the mooring load components. The final verification examines the time-domain signals directly, complemented by Welch-smoothed power spectral density (PSD) comparisons. Figures 6.34 and 6.35 present the true and predicted horizontal and vertical load compo-

nents, respectively, for three example samples from the dataset. Both time-domain





Figure 6.34: True versus predicted mooring load component F_x .

and frequency-domain comparisons show that the surrogate accurately reproduces the dynamic behaviour of the mooring loads. For the low-frequency dominated case (sample 129), the predicted signals closely follow the true signals, with minimal phase lag and amplitude deviation. The corresponding PSDs demonstrate excellent alignment across the dominant frequency range. In the other cases (samples 2026 and 8888), characterised by a broader spectral content, the agreement remains good, although the surrogate tends to overpredict the range 0.02 Hz to 0.07 Hz, and underpredict the higher-frequency load. Overall, the surrogate model preserves both the mean trends and the dynamic characteristics of the mooring load components in diverse loading conditions.

Figure 6.36 compares the mooring line tension time series (left) and the corresponding rainflow cycle process (right) obtained using the mooring surrogate-augmented model (C3) and OpenFAST for the original design under the rated wind speed condition. As previously observed for the samples within the dataset, the surrogate reproduces the general temporal trends from OpenFAST reasonably well. However, some underestimation of high-frequency tension fluctuations is noted. As shown in the rain-


 Figure 6.35: True versus predicted mooring load component F_z .

flow diagram, these small high-frequency discrepancies result in fewer low-amplitude cycles being counted. Moreover, the peak tensions are underestimated, despite the overestimation of the peak motions in the underlying motion solution (see Figure 6.11, Section 6.3.2). These effects together result in an overall underestimation of cumulative fatigue damage.




Figure 6.36: Tension and stress cycles counting for the mooring line (fairlead), comparison between OpenFAST and the surrogate model. Case 3.2 (rated conditions).

It should be noted that these findings pertain to a specific design under a specific environmental condition; a perfect match is neither expected nor required. The surrogate model is designed to provide accurate predictions on average across the design space, as demonstrated earlier in this section, and is sufficient for its intended use in

optimisation.

Notably, the implementation of the surrogate models within the coupled framework enabled the level of simulation efficiency required for the design optimisation studies. Table 6.15 summarises the average computational time for a single design evaluation, including 12 fatigue and one ultimate strength environmental condition. All simulations were executed on a single CPU (Intel[®] Xeon[®] Silver 4216 CPU at 2.10 GHz). The results highlight the substantial reduction in computational cost achieved by the surrogate-augmented model (C3) compared with the physics-based models (C1 and C2). Application of the BEM surrogate halved the solution time of the hybrid frequency-time domain model (C2). The application of the mooring surrogate model further reduced the simulation time by a factor of two. Overall, the surrogate augmented model only takes about 5 min to provide a comprehensive evaluation of the FOWT performance in 13 loading conditions.

Table 6.15: Average simulation time per design evaluation for models C1–C3 under 12 fatigue and 1 ultimate strength condition.

Model	BEM model	Mooring model	Computation time (s)
C1	Capytaine	Quasi-static	923.6
C2	Capytaine	MoorPy on motion time series	1232.2
C3	Surrogate	MoorPy on motion time series	586.8
C3	Surrogate	Surrogate	289.9

6.5 Case study

To demonstrate the impact of the modelling approaches, a series of support-structure optimisation exercises is presented, each employing one of the three models: C1 (physics-based frequency-domain), C2 (physics-based hybrid-domain), and C3 (surrogate-augmented hybrid-domain). Rather than demonstrating a fully realistic design workflow, the aim of this exercise is to isolate and evaluate the influence of the modelling approach on the optimisation process and its outcomes. For this purpose, a deliberately minimal optimisation problem is formulated.

The objective is to minimise the material cost of the mooring lines¹² by varying the mooring line shape parameter (α_{moor}), chain diameter (d_{moor}), anchor radius (r_{anchor}), and the fairlead vertical position normalised by the draft (z_F/T),¹³ subject to the constraints of the maximum offset and dynamic tension, all evaluated in the rated-wind speed condition. The line-shape parameter varies between 0 and 1, corresponding respectively to the shortest and longest feasible line lengths for the given anchor and fairlead positions. The constraint limits are set to those of the reference design (first iteration), ensuring that any cost reduction is accompanied by maintained or improved overall performance. The optimisation problems are solved using the IPOPT optimiser [283] with finite-difference derivatives.

As shown in Figure 6.37, all optimisation cases—despite following different optimisation paths—ultimately produced significantly more economical designs. In every case, the optimisation converged to the lowest allowable values of α_{moor} and the chain diameter, relying primarily on adjustments to the fairlead and anchor positions as the main cost levers. The two physics-based models (C1 and C2) arrived at nearly identical final designs, differing only slightly in the anchor radius. In contrast, the optimisation using model C3 converged to a distinct and cheaper solution, characterised by a smaller anchor radius, a lower fairlead position, and a shorter line length.

The dynamic tension constraint was active at the final solution in all three optimisation cases, indicating that it governed the optimal designs. However, the optimisation trajectories differed significantly between the physics-based models (C1, C2) and the surrogate-augmented model (C3).

For C1 and C2, the optimiser initially prioritised objective reduction, moving into the infeasible region with respect to the dynamic tension constraint. As the optimisation progressed, the constraint violation penalties became increasingly influential (consistent with IPOPT’s interior-point strategy), steering the solution back toward the feasibility boundary and increasing the cost slightly. Consequently, the final designs were reached by approaching the constraint from the infeasible side, with objective

¹²The values presented refer to the cost of a single line.

¹³The design variables are defined and illustrated in Section 4.2.




Figure 6.37: Comparison of optimisation process with models C1, C2, and C3. The dashed curve indicates the original design.

improvements ultimately limited by the need to satisfy the constraint.

In contrast, for C3, after an initial moderate reduction in objective value, further improvements occurred while remaining on the feasible side of the constraint boundary. The optimiser approached the limiting tension value from within the feasible region, allowing a smoother search trajectory and more effective exploration of feasible descent directions. This behaviour is reflected in the markedly smoother optimisation history.

The differing trajectories can be attributed to discrepancies in the dynamic tension and tension gradients predictions between the models. Since gradient information determines the search direction at each iteration, small differences in derivative estimates can lead to substantially different optimisation paths. In a non-convex design space with multiple local minima, such path dependency is expected. If the problem were convex, all optimisations would converge to the same solution regardless of initial search direction; however, the presence of multiple local minima makes the final design sensitive to the early gradient information and subsequent trajectory. In the present case, the surrogate-augmented optimisation converged to a solution with a lower objective value than that obtained using the physics-based models. However, this outcome should not be interpreted as systematically superior performance. Rather, in a non-

convex setting, differences in gradient information are expected to yield different local optima, which may be better or worse depending on the specific landscape encountered during the search.

Note that because no constraint was imposed on the direction of the load at the anchors,¹⁴ all optimal designs resulted in mooring lines that do not rest on the seabed in the equilibrium position, which is consistent with the objective of minimising cost. In practice, however, the savings associated with shorter lines could be offset by higher anchor costs, as traditional drag-embedded anchors might no longer be suitable for such configurations.

The initial design and the three optimised designs were subsequently analysed in OpenFAST to verify the improvements using a single consistent model. As summarised in Table 6.16, all optimisations produced stiffer mooring systems, achieving more than a 50 % reduction in maximum offset, albeit with an increase in pretension of approximately 30 %. When evaluated with the higher-fidelity model, the optimised configurations also exhibited higher dynamic tensions: around a 10 % increase for the surrogate-based optimisation and 25 % and 45 % increases for the two physics-based optimisations. This occurred despite the tension standard deviation constraint being active in the final iteration, indicating that modelling inaccuracies influenced the optimisation outcomes. The surrogate model resulted in the smallest increase in dynamic tension in the final design. Importantly, while all optimisations were relatively fast, the surrogate model delivered the final result in only 15 min, making it a valuable tool for supporting conceptual design with high potential for solving more complex optimisation problems.

6.6 Conclusion

This section provides a summary of the chapter and outlines the specific contributions to knowledge made.

¹⁴Purposefully, to keep the problem minimal as explained at the beginning of the section.

Table 6.16: Initial and optimised designs comparison and verification. Values based on OpenFAST simulation results. Tension relates to the upwind line tension.

	Initial design	C1 best	C2 best	C3 best
Mooring line cost (\$)	227 410	139 745	127 140	115 281
Mean + standard dev. offset (m)	33.55	9.64	12.84	13.23
Mean tension (kN)	1200	1660	1600	1560
Standard dev. of tension (kN)	70.35	88.7	101.0	76.3
Total optimisation time (min)	N/A	26.9	106.9	15.7

6.6.1 Summary

This Chapter addressed the need for an efficient coupled dynamic modelling framework suitable for the conceptual design optimisation of FOWT support structures. The goal was to balance computational efficiency with sufficient accuracy and fidelity to enable realistic design solutions. The developed model evaluates key performance metrics across ultimate, fatigue, and serviceability limit states, supporting a holistic and integrated design approach.

The coupled model was formulated by first establishing a fully converged static equilibrium solution, followed by dynamic simulations around this equilibrium. External loads and system responses were then mapped to internal loads in the tower and mooring lines. Increasingly complex model variants were developed and systematically verified against a fully coupled nonlinear reference.

The baseline model relied entirely on frequency-domain solutions for motion and load responses. However, it exhibited significant deficiencies under large mean offsets and high sea states, failing to capture the nonlinear and wave-frequency components of mooring tension and tower base bending moments with sufficient accuracy. As a result, the fatigue analysis was unreliable. The spectral-based fatigue assessment was also inadequate, as the assumption of a narrow-band stress process was violated under combined wind-wave loading. Additionally, the approach was computationally inefficient due to the reliance on boundary element methods for solving potential flow problems. These limitations are consistent with findings reported in the existing literature.

The next model introduced a transformation of frequency-domain motions into the

time domain, enabling direct calculation of structural loads and partial inclusion of motion-load nonlinearities. This significantly improved the accuracy of the tower loads prediction, with only a negligible increase in computation time. However, implementing either a quasi-static or dynamic lumped-mass mooring model in the time domain introduced a new efficiency bottleneck, yet substantial errors persisted in predicting mooring tensions, which propagated into the fatigue analysis. Although this model employed the rainflow counting method for fatigue estimation, inaccuracies in stress prediction continued to compromise the reliability of the fatigue analysis.

These limitations were addressed in the next model through the development of a data-driven machine learning surrogate for mooring loads in the time domain. This approach represents a significant advancement over state-of-the-art methods, as its dataset construction strategy is agnostic to both environmental conditions and the design of the floating platform. It relies solely on fairlead motion time series and mooring design parameters, making it broadly applicable across different projects while delivering near-instantaneous predictions with accuracy comparable to time-domain dynamic models. A separate surrogate model was also developed to overcome the efficiency bottleneck associated with the boundary element method solver.

The case study demonstrated that all three modelling approaches enabled substantial reductions in mooring line cost. The two physics-based models (C1 and C2) converged to similar solutions, whereas the surrogate-augmented model (C3) identified a distinct and notably cheaper configuration. This difference was attributed to variations in gradient information, which led to different optimisation trajectories within the non-convex design space. When evaluated in OpenFAST, all optimised designs showed increased stiffness and reduced offset, but also higher dynamic tensions, reflecting the influence of modelling inaccuracies on the active constraints. The surrogate-augmented optimisation produced the design whose dynamic-tension predictions aligned most closely with OpenFAST results. It delivered the solution in only 15 min, substantially faster than the physics-based models. Overall, these findings highlight the potential of the surrogate-augmented approach to support early-stage design by enabling rapid and reliable exploration of the design space.

6.6.2 Specific contributions

The following contributions to knowledge were made:

- Development of a hybrid frequency-time domain coupled model of dynamics for evaluation of design metrics, including FOWT motions, as well as ultimate and fatigue tower and mooring loads.
- Development of a database for the hydrodynamic coefficients surrogate model training and implementation of the model in the optimisation framework.
- Development of a database for the mooring tension and fatigue surrogate model agnostic to environmental conditions and floating platform design, and implementation of the model in the optimisation framework.

Chapter 7

Optimisation case studies

7.1 Introduction

These case studies aim to demonstrate how the findings and models developed in the previous chapters can be applied in practical design optimisation scenarios. Specifically, it seeks to address four key questions:

1. Does a fully integrated design of the tower, floating platform, and mooring system yield a more balanced and higher-performing floating offshore wind turbine (FOWT) compared with the traditional sequential design process?
2. How important is the inclusion of fatigue constraints for the tower and the mooring system at the conceptual design stage?
3. What is the impact of the new natural-frequency constraint formulation on the optimisation process and the resulting designs?
4. Is the instantaneous centre of rotation (ICR) a useful design metric?

In addressing these questions, the study further demonstrates that the proposed hybrid-domain, hybrid-paradigm modelling framework—together with the reformulated design problem—enables integrated optimisation to be performed efficiently and with appropriate realism, all within practical computational budgets at the conceptual design stage.

This study employs the gradient-based optimisation algorithm IPOPT [283]. IPOPT solves nonlinear optimisation problems of the form

$$\min_x f(x) \quad (7.1)$$

$$\text{s.t. } g_L \leq g(x) \leq g_U, \quad (7.2)$$

$$x_L \leq x \leq x_U, \quad (7.3)$$

where x is the vector of design variables, $f(x)$ is the design objective, $g(x)$ are the design constraints. IPOPT is a primal–dual interior-point method, in which the design variables, Lagrange multipliers associated with the constraints and bounds, and the slack variables introduced by the interior-point formulation are solved for simultaneously. Inequality constraints are handled by introducing slack variables (s) and a logarithmic barrier term ($-\mu \sum_i \ln(s_i)$) to enforce strict feasibility. IPOPT therefore solves a sequence of barrier subproblems

$$\min_{x,s} \phi_\mu(x, s) = f(x) - \mu \sum_i \ln(s_i), \quad (7.4)$$

subject to $g(x) + s = 0$, where $\mu > 0$ is the barrier parameter. At each iteration, a Newton step of the perturbed KKT¹ system is computed, and a line search ensures sufficient descent while maintaining $s_i > 0$. As $\mu \rightarrow 0$, the algorithm approaches a solution of the original constrained problem. Solving the Newton system requires first-order derivatives of the objective and constraint functions (an illustrative analytical example is provided in Appendix K).

Although analytical gradients are typically preferred for computational efficiency, the hybrid-domain modelling framework developed in this thesis is sufficiently fast to allow numerical finite-difference derivatives to be evaluated within practical runtimes. Furthermore, several components of the model are not differentiable, which remains a recognised bottleneck in the field; ongoing research efforts continue to address this limitation (e.g. [57]).

¹Karush–Kuhn–Tucker conditions - see Appendix K for details.

The optimisation approach used here is local and no multi-start strategy is applied. Consequently, the “optimal” designs reported in this chapter should be interpreted as locally optimal or improved designs. In addition, the inclusion of surrogate models introduces noise into the design evaluations, meaning that strict satisfaction of the optimality criteria (KKT conditions) cannot be guaranteed. Instead, the optimisation runs are terminated manually once all design variables and functions have converged sufficiently, and all constraint violations are acceptably small in engineering terms.

IPOPT is known to be sensitive to poor problem scaling: when the functions of interest vary over several orders of magnitude, this may lead to very slow convergence or no convergence at all. Therefore, all objective and constraint functions are scaled to lie within the 0–1 range. Commonly, such scaling is derived solely from function values at the first optimisation iteration. However, this approach may lead to poorly conditioned scaling elsewhere in the design space. To mitigate this risk, each optimisation is executed twice. The first run is terminated after a limited number of iterations (10–20) and is used exclusively to estimate appropriate scaling factors for all objectives and constraints. By basing the scaling on several early iterations rather than on the first iteration alone, the resulting scaling better reflects early optimisation dynamics and is less sensitive to the initial design point. These scaling factors are then applied in the second run to improve the numerical performance of IPOPT.

The C3 model introduced in Chapter 3.2.1 is used throughout this study. The design variables, objectives, constraints, and environmental conditions are defined in Chapter 2.2.1. For clarity and completeness, these elements are briefly reiterated for each test case in the subsequent sections.

The purpose of these exercises is to answer the four specific questions posed at the beginning of this chapter, rather than to demonstrate a real design process in its full complexity. The results of the three tests performed to address these questions are discussed in the following sections.




Figure 7.1: Test 1 workflow: sequential system optimisation.

7.2 Test 1: sequential versus integrated design

In this test, two series of optimisations are performed: a sequential design of the tower, floating platform, and mooring system, and an integrated design of all three subsystems. While both approaches are subject to the same set of constraints,² they differ in the formulation of the design objective and the selection of design variables. As illustrated in Figure 7.1, the sequential approach restricts the objective and variables to those of the subsystem currently being designed, whereas the integrated approach considers the full set of system-level variables and a system-level objective.

Each optimisation aims to minimise the material cost of the relevant subsystem (tower, platform, or mooring lines), subject to four dynamic-behaviour constraints (pitch stiffness, horizontal offset, tower tilt, and nacelle acceleration) and seven structural-integrity constraints (tower shell buckling and yielding, platform column buckling and yielding, mooring line strength, and vertical load at the anchors). The motion-related constraints and the vertical load on the anchors constraint are evaluated at the rated wind speed, while the strength utilisation at the tower base and in the mooring lines is evaluated under two conditions: high thrust and high significant wave height, as described in Section 4.4.2.

²All constraints are retained to ensure direct comparability between cases and to avoid subjective assignment of constraints to subproblems. This slightly (but acceptably) increases solution complexity due to the larger Jacobian handled by IPOPT.

The design variables differ according to the subsystem being optimised. For the tower, the variables are the tower-base radius and the ratio of tower-base wall thickness to diameter. This deliberately minimal set provides effective control over the tower's mass, stiffness, and structural strength, while keeping the problem dimensionality low to focus on interactions between subsystems and optimisation behaviour rather than on detailed structural design. For the platform, the design variables include the draft and a vector of five spar radii (a transition piece is introduced between the waterline and 10 m above the waterline, allowing the platform top and tower base to have different diameters and wall thicknesses). For the mooring system, the variables are the line diameter, the line-shape parameter, the fairlead vertical position relative to the platform draft, and the anchor radius. All variables are illustrated in Figures 4.1 and 4.3 in Chapter 4.

Note that the order of the sequential subsystem optimisations may influence the final results. The order adopted here—tower, platform, and mooring—represents one of several possible sequencing strategies and reflects common engineering practice at the global sizing or conceptual design stage.

Tower optimisation As illustrated in the workflow diagram in Figure 7.1, the designs of the platform and mooring system are initially assumed to be identical to the original OC3 configuration. The optimisation history for the tower is shown in Figure 7.2. Only the constraints that become active at any stage of the optimisation are displayed, for clarity, and the curves are coloured by iteration number for easier reference between the subplots.

The design of the tower was driven by the tower shell buckling utilisation limit of 1.16,³ with the remaining constraints well within the feasible ranges. A moderate reduction in the tower material cost of 12% was realised by a small decrease in the radius accompanied by a relatively small increase in the compactness ratio (i.e., decrease in the wall thickness) – exact values are presented in Table 7.1.

³The buckling utilisation limit of 1.0 typically denotes the onset of buckling; the limit adopted here corresponds to the utilisation predicted by the current model for the reference design.




Figure 7.2: Test 1: Tower optimisation history. Black dashed line – initial value.




Figure 7.3: Test 1: Floating platform optimisation history. Waterline radius – r_4 . Black dashed line – initial value, blue dashed line – constraint limiting value.

Platform optimisation As illustrated in Figure 7.1, the tower design from the previous step, and the original OC3 mooring design, are the starting point for the floating platform optimisation. The optimisation history is presented in Figure 7.3.

Three of the platform radii (r_1 , r_2 , r_3) are driven to their lower bounds, and r_0 (keel radius) decreases, all reducing platform mass and cost. The waterline radius r_4 , however, increases monotonically from its baseline of 3.25 m to approximately 6.5 m, compensating the loss of pitch restoring stiffness C_{44} that would otherwise result from the reduced mass and geometry. Increasing r_4 raises both z_G and z_B , with z_G rising faster due to more steel concentrated near the waterline and less ballast at the keel,

which weakens the stabilising gravity term $-Mgz_G$. However, the rising z_B reduces the destabilising buoyancy separation term ρgVz_B by a greater margin, and the waterplane moment ρgI_{wp} , while an order of magnitude smaller, contributes in the same direction. The net effect is a recovery of C_{44} sufficient to reduce pitch displacement and, consequently, the tower shell buckling utilisation which is active at convergence.

Arguably, increasing the keel radius r_0 could provide an alternative route to recovering pitch restoring stiffness; while it would increase the stabilising gravity term $-Mgz_G$ effectively, it would simultaneously lower z_B , amplifying the destabilising buoyancy separation term ρgVz_B . The trade-off between these two competing effects is non-trivial, and its resolution is left to the optimiser. It is possible that different starting points would lead the local optimiser to select either r_0 or r_4 as the primary stability lever in different realisations.

Consistent with the theoretical behaviour of the IPOT algorithm detailed in Section K.0.5 in Appendix K, the draft T decreases sharply in the first few iterations from its baseline of 120 m, because the barrier-based merit function is initially dominated by the objective term — reducing T being a very effective way of reducing cost, given that mass scales linearly with draft. After a few iterations, constraint satisfaction becomes increasingly influential, leading to a partial recovery in T and a slower reduction in cost, primarily driven by small adjustments in the spar radii. The final design achieved a substantial cost reduction of 23 %, with all constraints satisfied.

Mooring optimisation Following the logic of Figure 7.1, the tower and platform designs from the previous steps are the starting point for the moorings optimisation. The optimisation history is presented in Figure 7.4.

The optimiser converged smoothly towards a substantially lower-cost solution, achieving a 29 % reduction. Multiple constraints were active at the final point, including those related to the horizontal offset (the maximum offset and vertical load on the anchor), and the strength criteria at the tower and mooring system. Importantly, mooring yield utilisation in the high-wave condition (ULS_2) became the design driver, indicating that the optimiser approached the lightest possible design.




Figure 7.4: Test 1: Mooring line optimisation history. Black dashed line – initial value, blue dashed line – constraint limiting value.

Cost reduction was primarily driven by decreasing the chain diameter d , which reduces the submerged weight per unit length and hence both the fabrication cost and the mooring tensions T_h and T_{fair} . However, reducing d also reduces the linearised mooring stiffness $\partial T_h / \partial x_{\text{fair}}$, weakening the restoring force available to limit horizontal platform offset.

Constraint satisfaction was primarily achieved by increasing the fairlead-anchor distance and the suspended line length, the latter controlled through the line-shape parameter α_{moor} . Increasing the suspended line length reduces both tensions and stiffness, all else being equal, which has a negative impact on offset, but helps to lower the vertical anchor load. Increasing the fairlead-anchor distance raises tensions and stiffness, compensating for the stiffness lost through diameter reduction and line length increase. Together, these geometric adjustments balanced the horizontal offset, line ULS utilisation, and anchor load constraints, all of which were active at convergence.

Second pass through sequential optimisation A single pass through the sequential optimisation workflow shown in Figure 7.1 already produced a substantially more economical system. This approach, however, captures only one-way interactions between the tower, floating platform, and mooring system designs, suggesting that further improvement might still be possible.

As shown in Table 7.1, the second iteration of the workflow, initialised from the best designs obtained in the first pass, did not lead to any further meaningful improvement. The second tower optimisation failed to identify any feasible solutions with a lower cost. The second platform optimisation converged to a slightly more expensive design, primarily by reducing marginal constraint violations close to their numerical tolerances. Using this modified platform as the starting point, the second mooring optimisation achieved only a marginal, practically negligible, reduction in mooring cost. Overall, the second pass through the sequential optimisation process resulted in a total system's variables and cost comparable to that obtained after the first pass, indicating that the sequential approach was unable to escape the identified local optimum.

Integrated optimisation initialised from original design Next, an integrated optimisation of the tower, floating platform, and mooring system was performed, in which the design variables of all subsystems were varied simultaneously to identify the configuration with the lowest total material cost.

As illustrated in Figure 7.5, the optimiser converges to the minimum in fewer than 50 iterations – fewer than the total iterations of the 1st pass through sequential optimisations. This behaviour illustrates the efficiency of the gradient-based optimisation approach in a reasonably smooth and well-conditioned design space. Despite the increased dimensionality of the coupled problem, comprising 12 design variables, the optimiser exploits gradient information to efficiently navigate the design space and rapidly identify descent directions.

The platform adopts the same qualitative solution as in the sequential case: r_4 increases to recover pitch restoring stiffness C_{44} , while r_1 – r_3 collapse to their lower bounds and r_0 reduces, all driving mass and cost down. The pitch displacement constraint $\eta_{5,\max}$ is active at convergence, confirming that r_4 has been increased only as far as necessary to satisfy the stability margin.

The active constraints at convergence differ from the sequential result: the horizontal offset $\eta_{1,\max}/\text{depth}$, the vertical anchor load $F_{z,\text{moor}}$, and tower shell buckling are all active, while the mooring yield utilisation under the high-wave condition (ULS_2) is no longer design-driving. This reflects the greater freedom available in the integrated problem: with all subsystem variables adjusting simultaneously, the optimiser finds a different low-cost design where mooring strength is no longer the limiting factor. The mooring solution follows the same qualitative direction as the sequential case — reduced chain diameter d from 70 mm to approximately 58 mm, increased fairlead-anchor distance r_{anchor} from 845 m to approximately 925 m, and increased suspended line length via α_{moor} — but converges to a different balance point among the active constraints.

The updated tower features a base radius R_{twr} settling near 3.4 m and a D/t_{twr} ratio of approximately 260, above its initial value of 250, with tower shell buckling active at convergence. The wall thickness change remains marginal and indistinguishable in practical engineering or manufacturing terms.




Figure 7.5: Test 1: Integrated support structure optimisation history. Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4 .

Integrated optimisation initialised from best sequential optimisation design

The sequential optimisation failed to escape the identified local optimum. As a final experiment in Test 1, an integrated optimisation was therefore performed again, this time initialised from the locally optimal design obtained through the sequential approach. The corresponding optimisation history is shown in Figure 7.6.

The optimisation successfully escaped the local minimum and arrived at a substantially lower-cost solution of approximately 2700 kUSD, comparable to the result obtained from the original initialisation. However, in contrast to the smooth descent observed in that run, the early iterations here exhibit a pronounced transient: the objective spikes to approximately 3400 kUSD and multiple constraints spike briefly reaching their limits before retreating, reflecting the larger perturbation required to move away from the sequential local minimum. Convergence requires approximately 70–80 iterations, roughly double that of the first integrated run.

The converged design mostly follows the same qualitative trends and the same set of active constraints as the previous integrated run. The draft, however, settles near 113 m, somewhat lower than the 119 m obtained from the original initialisation, with the associated reduction in pitch restoring stiffness compensated by a correspondingly larger waterline radius r_4 , which converges to a slightly higher value than in the first integrated run.

As summarised in Table 7.1, both integrated optimisations substantially outperform the sequential result. The close agreement in final cost and active constraint set between two runs initialised from markedly different starting points increases confidence in the robustness of the optimisation and suggests that the identified solution represents a near-global optimum for Test 1.

Summary of Test 1 The results of all optimisations are compared in Table 7.1 and Figure 7.7. The sequential (light green bar) and integrated (grey bar) optimisations resulted in 12% and 22% total cost reduction, respectively. Compared with the sequential approach, the integrated optimisation produced a more expensive mooring system. However, this additional cost enabled a significant reduction in platform cost,

Chapter 7. Optimisation case studies




Figure 7.6: Test 1: Integrated support structure optimisation history initialised from the end point of the sequential optimisation. Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4 .

ultimately yielding a cheaper overall system. Such a trade-off could not be achieved with the sequential approach, which inherently selects the least expensive option for each subsystem independently.

The physical mechanism underlying this trade-off can be understood by writing the

Table 7.1: Test 1: Comparison between the initial FOWT design and the designs optimised in a sequential and integrated manner. Costs are normalised relative to the initial total cost of 3,452,500 USD.

	Unit	Initial design	Seq. opt. (1st pass)	Seq. opt. (2nd pass)	Int. opt. (from seq. 2)	Int. opt. (from initial)
r_{twr}	m	3.25	3.08	3.04	3.40	3.48
D/t	–	240.7	250.7	239.6	261.8	271.6
(t_{twr})	m	0.0270	0.0246	0.0254	0.0261	0.0256
T	m	120.00	120.31	120.88	113.60	118.97
r_1 (keel)	m	4.70	3.95	4.10	3.32	3.27
r_2	m	4.70	3.00	3.19	3.03	3.00
r_3	m	4.70	3.00	3.11	3.01	3.00
r_4	m	3.98	3.00	3.11	3.06	3.00
r_5 (WL)	m	3.25	6.80	6.85	5.69	5.28
α_{moor}	–	0.121	0.116	0.117	0.116	0.117
d	mm	70.0	59.34	58.89	58.99	58.76
$z_{fairlead}/T$	–	0.58	0.59	0.60	0.22	0.20
R_{anchor}	m	845.2	854.1	855.7	922.6	928.7
Material cost (normalised to initial total = 1.0)						
Tower	–	0.18	0.16	0.16	0.16	0.18
Platform	–	0.62	0.56	0.59	0.54	0.45
Mooring lines $\times 3$	–	0.20	0.13	0.13	0.08	0.15
Total	–	1.00	0.85	0.88	0.78	0.79




Figure 7.7: Test 1: Comparison between the initial and optimised FOWT designs.

pitch equilibrium equation about the mooring attachment point:

$$[C_h + C_g] \theta = F_{aero} |z_{hub} - z_{moor}| \quad (7.5)$$

where C_h and C_g are the hydrostatic and gravitational restoring coefficients defined

about the fairlead, θ is the pitch angle, F_{aero} is the aerodynamic thrust, and z_{hub} and z_{moor} are the heights of the hub and mooring attachment point, respectively. By raising the fairlead attachment point z_{moor} , the integrated optimiser reduces the overturning moment arm $|z_{\text{hub}} - z_{\text{moor}}|$, so that a smaller combined restoring coefficient $[C_h + C_g]$ is sufficient to maintain the same equilibrium pitch angle θ . Since reducing displaced volume is the most cost-effective way to reduce C_h , the platform radii can be decreased, yielding a lighter and cheaper platform. The mooring system incurs a higher cost to achieve the elevated fairlead position, but the resulting platform savings more than compensate, producing a lower total system cost. This cross-subsystem trade-off is only accessible to the integrated formulation, where the fairlead position and platform geometry are optimised simultaneously.

A potential alternative would be a nested optimisation architecture (for example, following one of the distributed multidisciplinary design analysis and optimisation, MDAO, formulations), with an outer loop minimising the total system cost. However, given the simplicity, interpretability, and low computational demand of the direct integrated approach for FOWT design, adopting a nested formulation is unlikely to provide substantial additional value for the problem considered in this work. Based on the results of this test case, there is limited evidence to suggest that further decomposition of the optimisation problem would be beneficial, which supports the use of monolithic MDAO architectures unless the design problem is substantially redefined.

7.3 Test 2: the impact of fatigue constraints

The optimum design obtained from the integrated optimisation satisfied all constraints considered up to this point; however, fatigue damage had not yet been included in the formulation. A subsequent evaluation with model C3 revealed that the tower-base fatigue damage substantially exceeds the permissible limit of 1, reaching a value of 26.9.⁴ This clearly indicates that fatigue must be explicitly addressed in the design process

⁴Note that high fatigue accumulation results are not uncommon, as fatigue is very sensitive to underlying assumptions and stress evaluation accuracy, values in the order of 20 in otherwise optimised designs were presented for example in [284].

to obtain a structurally viable solution. In the next step, fatigue damage for both the tower–base section and the mooring line at the first chain link at the fairlead is incorporated into the integrated optimisation of the support structure. All other elements of the formulation are kept unchanged, ensuring a direct assessment of the isolated effect of fatigue constraints.

As seen in the optimisation convergence history in Figure 7.8, the introduction of the tower and mooring fatigue constraints led the algorithm to converge to a higher-cost design with significantly reduced fatigue levels, as expected. The reduction in tower fatigue was achieved primarily through a decrease in tower loading, driven by a reduction in pitch motion ($\eta_{5,\max}$). This was mainly accomplished by changing the distribution of platform radii: a substantial increase in the platform radii near the keel increased the system inertia and lowered the centre of ballast, increasing the pitch stiffness; meanwhile, radii near the waterline in the wave-affected zone were reduced, lowering wave excitation and further contributing to the reduction in platform motions and fatigue damage. While the tower geometry itself changed only marginally, the optimisation shifted toward a shorter-line mooring system. This adjustment helped to lower the total cost, and the mooring system no longer needed to contribute as much to motion suppression, since the heavier and more inertial platform already provided improved motion performance.

Increasing the tower base radius and/or wall thickness could, in principle, provide an effective means of controlling fatigue damage by increasing the structural strength. However, while such modifications enhance strength, a bulkier tower also shifts the system’s centre of mass upward. This results in increased pitch motions and, consequently, higher dynamic loads at the tower base, which can have a detrimental effect on fatigue performance. Assessing this trade-off within a manual design process is challenging, as the competing effects are strongly coupled and non-intuitive. In contrast, the optimisation framework enabled the systematic identification of a more efficient design pathway that reduces overall system cost while maintaining adequate structural reliability.

To further support this interpretation, the optimisation was repeated using a deliberately stocky tower as the initial design. As shown in Figure 7.9, the optimiser

Chapter 7. Optimisation case studies




Figure 7.8: Test 2: Support structure optimisation history with fatigue constraints (1st start point). Black dashed line – initial value, blue dashed line – constraint limiting value. Waterline radius – r_4 .

initially explored the increased-strength design direction (high R_{twr} and low D/t) but ultimately converged away from this solution, indicating that this local optimum was less favourable than the alternative fatigue reduction pathway.

The final locally optimal designs are compared with the design obtained from the




Figure 7.9: Test 2: Support structure optimisation history with fatigue constraints (2nd start point). Black dashed line – initial value, blue dashed line – constraint limiting value.

optimisation without fatigue constraints in Table 7.2 and Figure 7.10. The previously discussed interdependence between the platform and mooring system is clearly reflected in the cost breakdown. Incorporating fatigue constraints increased the floating platform cost, more markedly for the first starting-point case, while simultaneously reducing the mooring system cost. As before, the tower cost remained largely unchanged. Accounting for fatigue increased the total system cost by approximately 8 % for the first starting-point case; nevertheless, the resulting design remains about 15 % cheaper than the original baseline configuration. The second starting-point optimisation achieved fatigue reduction with no net increase in total cost, more effectively balancing the increased tower base strength and platform inertia against the reduced mooring system cost.

Table 7.2: Test 2: Comparison between the best FOWT designs from optimisations with and without fatigue constraints. Costs are normalised relative to the cost of 2,722,703 USD.

	Unit	Integrated opt. (no fatigue)	Integrated opt. (with fatigue) (1st start)	Integrated opt. (with fatigue) (2nd start)
r_{twr}	m	3.48	3.44	3.52
D/t	–	271.6	266.54	277.96
(t_{twr})	m	0.0256	0.0258	0.0253
T	m	118.97	112.71	119.97
r_1 (<i>keel</i>)	m	3.27	4.17	3.58
r_2	m	3.00	3.46	3.00
r_3	m	3.00	3.20	3.00
r_4	m	3.00	4.67	3.88
r_5 (<i>WL</i>)	m	5.28	3.32	3.91
α_{moor}	–	0.117	0.109	0.116
d	mm	58.76	61.17	58.59
$z_{fairlead}/T$	–	0.20	0.33	0.21
R_{anchor}	m	928.66	880.99	928.31
Material cost (normalised to initial total = 1.0)				
Tower	–	0.23	0.23	0.23
Platform	–	0.58	0.67	0.59
Mooring line $\times 3$	–	0.20	0.18	0.18
Total	–	1.00	1.08	1.00




Figure 7.10: Test 2: Comparison between the FOWT designs optimised with and without the fatigue constraints.

7.4 Test 3: the impact of natural frequency constraints

As anticipated from the frequency–penalty landscape in Figure 4.10 (Section 4.4.1), the resulting design space is highly nonconvex. Consequently, the local optimisation is repeated from multiple starting points. The resulting designs are summarised in Table 7.3 and Figures 7.11 and 7.12.




Figure 7.11: Test 3: Comparison of FOWT designs optimised with the natural frequency constraints locally from different starting points.

The dominant penalty terms at the optima are associated with the overlap between the three lowest natural frequencies and the wind excitation range, namely $P(f_{nat1,wind})$, $P(f_{nat2,wind})$, and $P(f_{nat3,wind})$, as well as the overlap between the fourth lowest natural frequency and the rotor 3P excitation range, $P(f_{nat4,3P})$. The resulting

Table 7.3: Test 3: Comparison of FOWT designs optimised with the natural frequency constraints locally from different starting points. Costs are normalised relative to the initial total cost of 3,452,500 USD.

	Unit	Initial	Opt. 1	Opt. 2	Opt. 3
<i>Design variables</i>					
R_{twr}	m	3.25	3.21	4.43	4.62
D/t_{twr}	–	240.7	240.1	245.2	213.3
(t_{twr})	m	0.0270	0.0267	0.0361	0.0433
T	m	120.0	119.5	119.6	115.4
α_{moor}	–	0.127	0.124	0.123	0.0
d	mm	70.0	69.89	64.95	71.27
r_{anchor}	m	853.1	837.8	844.7	759.7
z_{fairlead}/T	–	0.585	0.586	0.584	0.586
r_1	m	4.70	4.69	5.18	5.59
r_2	m	4.70	4.67	4.67	5.12
r_3	m	4.70	4.76	4.67	5.11
r_4	m	3.95	2.94	3.94	5.25
r_5	m	3.25	3.26	4.49	5.32
<i>Constraints (active)</i>					
$P(f_{\text{nat}_1, \text{wind}})$		0.22	0.21	0.24	0.08
$P(f_{\text{nat}_2, \text{wind}})$		0.04	0.07	0.04	0.03
$P(f_{\text{nat}_3, \text{wind}})$		0.04	0.03	0.02	0.02
$P(f_{\text{nat}_4, 3P})$		0.19	0.00	0.00	0.00
$\sum$		0.49	0.31	0.30	0.13
Material cost (normalised to initial total = 1.0)					
Tower		0.18	0.33	0.30	0.37
Platform		0.62	0.66	0.79	1.01
Mooring line $\times 3$		0.20	0.18	0.16	0.17
$\sum$		1.00	1.19	1.25	1.55




Figure 7.12: Test 3: Natural frequency penalty map for the FOWTs optimised locally from different starting points.

penalty landscape and the locations of the natural frequencies for the optimised designs are illustrated in Figure 7.12.

As summarised in Table 7.3 and Figure 7.11, none of the locally optimised solutions achieves a strictly zero value for all resonance-related penalty terms. Nevertheless, all optimised designs exhibit a substantial reduction in the total penalty compared to the initial configuration, with the summed value decreasing from 0.49 to between 0.31 and 0.13. This demonstrates that the proposed smooth penalty formulation systematically guides the optimisation away from resonance-prone regions.

A consistent outcome across all starting points is the complete elimination of the 3P-related overlap. The initial design exhibits a significant penalty associated with the overlap between the tower’s first fore–aft natural frequency and the rotor 3P excitation range, $P(f_{\text{nat}_4,3P})=0.19$, whereas this term is reduced to zero for all locally optimised solutions. Figure 7.12 shows that the initial configuration lies on the low-frequency side of a bimodal 3P penalty landscape, while all optimisation runs successfully traverse the steeper of the two penalty peaks and converge to the high-frequency, stiff–stiff regime. This indicates that the gradient-based optimisation process was able to traverse a nonconvex resonance landscape rather than becoming trapped in the nearest local minimum.

The residual penalty at the optima is dominated by the overlap between the system’s lowest natural frequency, associated with rigid-body surge motion, and the wind excitation range. For all locally optimised solutions, $P(f_{\text{nat}_1, \text{wind}})$ remains non-zero, indicating that complete avoidance of this overlap is not achievable within the explored design bounds and coupled system constraints. One optimisation run yields a markedly lower value of this term compared to the others, reflecting design changes that shift the stiffness–mass balance towards higher natural frequencies (primarily through shorter, larger-diameter mooring lines). However, this reduction is accompanied by a substantial increase in structural cost.

The comparison of the locally optimised solutions in Figure 7.11 reveals a clear trade-off between structural cost and resonance-related risk. While moderate reductions in total penalty can be achieved with relatively modest cost increases, the solution exhibiting the lowest total penalty corresponds to a substantially higher normalised material cost. Notably, this solution is characterised by a qualitatively different set of design variables rather than a refinement of the initial configuration, indicating a transition to a distinct design regime (region in the design space). This observation highlights the strongly nonconvex nature of the optimisation landscape and suggests that significant reductions in resonance exposure require more fundamental changes to the structural concept rather than incremental tuning.

Furthermore, Figure 7.12 shows that the tower natural frequency is shifted well beyond the minimum value required to achieve zero 3P penalty. This behaviour indicates nontrivial coupling between the resonance-related penalty terms: design modifications intended to mitigate one excitation mechanism simultaneously reshape the feasible region associated with others. The optimiser negotiates a coupled, multi-dimensional trade-off between competing metrics.

Overall, embedding prior knowledge of excitation likelihood and severity into smooth, differentiable penalty functions enabled the optimisation to produce a set of physically meaningful local optima spanning different cost–risk compromises. Rather than yielding a single binary-feasible solution, the proposed approach provided valuable decision-support information by revealing alternative design regimes and their associated trade-

offs, allowing the developer to select a configuration based on acceptable resonance exposure and structural cost.

Note that operating near or at resonance can be particularly detrimental in terms of fatigue damage, especially for lightly damped systems, as even moderate excitation levels may lead to significantly amplified cyclic loads. Fatigue constraints were, however, not included in this test, as the intent was to isolate and assess the effect of the proposed resonance penalty functions independently of other design considerations.

7.5 Test 4: ICR-informed design

Chapter 5 proposed an example formulation of the design problem formulation based on an unconventional metric: the instantaneous centre of rotation (ICR). In this example, the suitability of the ICR for informing the design of a spar support structure is investigated by direct inspection of the design space.

Using the ICR statistics and responses of interest identified as correlated in Chapter 5, Figure 7.13 illustrates the design space in the form of $2D$ trends (full-factorial design) for all combinations of the design variables selected in Chapter 5: the vertical centre of mass and fairlead location.

While the sensitivity of the mean x_{ICR} to each variable is largely independent of the other variable (top-left subplot), more interdependence is seen in the case of the standard deviation of z_{ICR} , where the rate of the ICR increase with the centre of mass depends on the fairlead position and vice versa. Regarding the loads, the mooring line tension at the fairlead is significantly more sensitive to variations in fairlead position than it is to the centre of mass. The impact of the two design variables on the tower base and blade root bending moments is more balanced, with both variables affecting the loads significantly.

As previously outlined in the diagram of Figure 5.30 and additionally illustrated in Figure 7.13, the design modifications considered in this case study follow the steps listed below:

- Lower the fairlead position to increase the mean x_{ICR} (top-left plot, red arrow).




Figure 7.13: Design space visualisation: ICR and response statistics as functions of two selected design variables for the spar-type FOWT.

- This leads to a reduction in the mean and dynamic mooring tension (top-middle and top-right plots, red arrows) and to an unfavourable increase in the tower base and blade root bending moments (bottom-middle and bottom-right plots, red arrows).
- Lower the platform centre of mass to lower the standard deviation of z_{ICR} (bottom-left plot, green arrow).
- This brings the tower base and blade root loads back close to the original values while retaining the reduced mooring tension.

The extent of the variable adjustments is limited by practical considerations: it may not be practical to construct the fairlead at the very bottom of the spar structure, and the centre of mass of the platform is limited by the ballast placement.⁵

To support reading the 2D maps, OpenFAST results of the original and the modified designs are compared in Tables 7.4–7.7 and Figure 7.14. As expected, the most significant change was seen in the mooring loads, with 26 % and 69 % reductions in the mean

⁵Implementing such adjustments may necessitate additional design modifications—for example, lowering the centre of mass could require a deeper draft structure. While some consideration has been given to the practicality of these design variations, their overall feasibility is not guaranteed.

and dynamic tension, respectively, improving both the ultimate strength and fatigue characteristics of the design. These changes were accompanied by shifting the mean of the body-fixed coordinate system ICR closer to the fairleads position. The softer mooring system resulted in vastly increased mean surge offset of 21.73 m. Although this was a 50 % increase over the original design, the value is still reasonable (7 % of the water depth assumed). Importantly, the reduction in mooring loads was achieved without significantly compromising the other structural loads: the only response negatively affected was the tower base bending moment, which increased by less than 5 %.

Table 7.4: Comparison of the original and the ICR-informed adjusted spar FOWT design responses. Body-fixed frame ICR values are in brackets.

Metric	Unit	Mean		Std	
		Original	Adjusted	Original	Adjusted
x_{ICR}	m	-2.13 (21.60)	-0.79 (12.63)	54.79 (54.80)	33.40 (33.35)
z_{ICR}	m	-28.69 (-23.27)	-24.31 (-28.86)	19.83 (19.58)	15.30 (15.36)
FAIRTEN2	N	1.09E+06	8.06E+05	4.24E+03	1.31E+03
TwrBsMyt	kN m	2.68E+04	2.64E+04	4.04E+03	4.21E+03
RootMyc1	kN m	3.83E+03	3.85E+03	2.37E+02	2.32E+02
Surge	m	13.93	21.73	0.12	0.13
Heave	m	-0.16	1.04	0.02	0.02
Pitch	°	1.62	1.53	0.06	0.06
NcIMUTAxS	m s^{-2}			0.11	0.11

Table 7.5: Comparison of the original and the ICR-informed adjusted spar FOWT design responses - absolute and relative percentage differences.

Metric	Unit	Mean		Std	
		Abs. diff	Rel. diff (%)	Abs. diff	Rel. diff (%)
x_{ICR}	m	1.34 (-8.97)	62.9 (-41.5)	-21.39 (-21.45)	-39.0 (-39.14)
z_{ICR}	m	4.38 (-5.59)	15.3 (-24.0)	-4.53 (-4.22)	-22.8 (-27.5)
FAIRTEN2	N	-2.84E+05	-26.1	-2.93E+03	-69.1
TwrBsMyt	kN m	-4.00E+02	-1.49	1.70E+02	4.21
RootMyc1	kN m	2.00E+01	0.52	-5.00	-2.11
Surge	m	7.80	56.0	0.01	8.33
Heave	m	1.20	750.0	0.00	0.00
Pitch	°	-0.09	-5.56	0.00	0.00
NcIMUTAxS	m s^{-2}	—	—	0.00	0.00


Although it was impossible to align the mean ICR with the fairleads using the design modifications considered, moving the ICR in that direction significantly reduced

Table 7.6: The original and the ICR-informed adjusted spar FOWT designs features.

Feature	Unit	Original	Adjusted
Fairlead z	m	-70.00	-110.00
Platform z_{CM}	m	-89.92	-100.00

Table 7.7: The original and the ICR-informed adjusted spar FOWT designs features - absolute and relative differences.

Feature	Unit	Abs. diff	Rel. diff (%)
Fairlead z	m	-40.00	-57.1
Platform z_{CM}	m	-10.08	-11.2

Figure 7.14: The responses of the modified spar FOWT design. Values normalised with the values of the original design as $1 + (x_{adjusted} - x_{original})/|x_{original}|$.

the mean and dynamic mooring loads. As already hinted at the end of Section 5.7, this adjustment had a negative impact on the tower and blade loads, which was mitigated by shifting the platform's centre of mass, thereby shifting the standard deviation of z_{ICR} and restoring these loads to their original levels. Even though the ICR is dynamic and exhibits significant variability, the case study showed that shifting the mean ICR towards specific locations can significantly reduce structural loads.

7.6 Conclusion

This section provides a summary of the chapter and outlines the specific contributions to knowledge made.

7.6.1 Summary

This chapter applied the models developed in Chapters 3.2.1 and 5 and the problem formulation introduced in Chapters 4 and 5 to perform design optimisation studies of the support structure, including the tower, floating platform, and mooring system.

The results demonstrated that the co-design of the floating platform and mooring system is important, as both significantly influence the dynamic behaviour of a FOWT. In particular, over-engineering of one subsystem can enable cost savings in the other, leading to a more cost-effective overall design. The case study showed that integrated optimisation of the FOWT system can be achieved at a computational cost lower than a series of sequential subsystem optimisations. Furthermore, fatigue assessment was shown to be a critical component of the design process. The additional cost required to reduce fatigue, achieved by increasing the system inertia and thereby lowering dynamic loads, was quantified. The proposed resonance constraint formulation was shown to function as an effective decision-support tool by embedding prior knowledge of excitation likelihood and severity into the penalty shapes, thereby providing informative risk gradients that enable meaningful trade-offs both between structural cost and resonance risk and among competing resonance mechanisms, beyond binary feasibility.

In the last case study on design informed by the instantaneous centre of rotation (ICR), the results demonstrated the effect of the ICR on the trade-off between dynamic loads at different locations on the structure. Shifting the ICR closer to the tower top reduced nacelle acceleration and tower base loads (which are driven by acceleration). In contrast, shifting the ICR closer to the fairleads reduced mooring line loads. Importantly, the study presented in this Chapter and a complementary study presented in [217] showed that the ICR should not be treated as an independent design variable. It can only be adjusted by modifying other design variables that also impact

the loads through mechanisms different from the ICR itself and, in some cases, may even counteract its effects. Because of these complexities, relying solely on the ICR to guide the design is insufficient. Instead, the ICR is useful for understanding one of the mechanisms that affect the global response and loads experienced by FOWTs.

7.6.2 Specific contributions

The following contributions to knowledge were made:

- An integrated optimisation of the FOWT support structure, enabling quantitative assessment of performance and cost trade-offs between the design of the tower, platform, and mooring system, with direct comparison of efficiency of the sequential and integrated optimisation approaches.
- Quantitative characterisation of the trade-off between structural fatigue life (tower and mooring system) and system cost within floating wind support structure optimisation, demonstrating the impact of fatigue constraints on optimal design solutions.
- Introduction of a probability- and power-aware natural-frequency constraint formulation that enables quantitative assessment of resonance risk within design optimisation.

Chapter 8

Conclusion and future work

This chapter presents the conclusions of the thesis and outlines directions for future research.

8.1 Context and problem

Floating offshore wind turbines (FOWTs) enable wind energy development in deeper waters with stronger and more consistent wind, increasing renewable generation and supporting climate targets. However, the combination of harsher offshore environments and complex subsystem interactions introduces major challenges for design, particularly in modelling and evaluating structural integrity across a range of failure mechanisms and conditions. Despite these complexities, current concept design practice still relies on simplified models and engineering judgment, with subsystems developed sequentially by different stakeholders. This slow, repetitive, and manual process is limited in scope and risks overlooking synergies or optimal solutions.

These limitations highlight the need for more integrated, computationally advanced design approaches. By treating the floating platform, mooring system, power export cable, tower, RNA, and controllers as a coupled system, it is possible to achieve a balanced trade-off between cost and performance, avoiding over-engineering of individual components. Casting the design as an optimisation problem allows numerical methods to efficiently explore large design spaces under realistic constraints, improving the like-

likelihood of finding globally optimal solutions at the conceptual stage, when changes are least costly to implement.

Design optimisation requires a practical evaluation approach that balances efficiency and accuracy by capturing the key design-driving effects at a reasonable computational cost. This requirement is further amplified by the iterative nature of optimisation and the long lifespan of FOWTs, which expose structures to a wide range of conditions — from moderate, recurring loads that drive fatigue to rare, extreme events imposing high-magnitude loads. Moreover, the stochastic nature of wind and wave excitation necessitates multiple random realisations, while the inherently multidisciplinary character of the problem — spanning aerodynamics, hydrodynamics, structural dynamics, electrical engineering, control, and systems engineering adds another layer of complexity.

Multidisciplinary design analysis and optimisation (MDAO) provides a structured methodology to address these challenges by integrating different domains into a unified framework of interconnected analysis blocks, coupled into a flexible, system-level model that simplifies development and supports collaboration. It ensures that interdependencies are properly captured, and facilitates the computation of total system derivatives, which are key to efficient, gradient-based optimisation that scales effectively with problem complexity.

Decomposing a model into components allows multi-fidelity approaches. Modules may differ in domain (time, simplified time, frequency, or hybrid time–frequency), in the number of degrees of freedom (from fully flexible to simplified rigid-body models), or in modelling paradigm (physics-based, data-driven, or physics-informed). These strategies can be combined to balance accuracy and efficiency across the coupled model. In particular, data-driven machine learning models trained on mid- or high-fidelity simulations offer near-instantaneous evaluations while retaining the accuracy of the underlying data.

The MDAO methodology requires formulating the optimisation problem in line with conceptual design goals and constraints, typically guided by offshore standards and regulations. At the same time, the early design phase allows flexibility and innovation,

where alternative metrics can complement or extend those prescribed by standards. One such metric is the instantaneous centre of rotation (ICR), which provides a compact characterisation of system behaviour. Incorporating the ICR into the design process can simplify the optimisation problem.

8.2 Aim and objectives

In this context, the thesis aims to improve the efficiency and scope of the FOWT support structure conceptual design optimisation through the application of the MDAO methodology and a multi-fidelity, multi-domain, physics-based surrogate-supported model of dynamics. As such, this research seeks to propose an optimisation problem formulation suitable for the concept design and appropriate design evaluation methods, thereby enabling a more informed and systematic approach to FOWT design from the earliest design stages.

To achieve this aim, the research was structured around three objectives: **O1**: Apply a multidisciplinary design analysis and optimisation (MDAO) framework to the integrated optimisation of the FOWT system, coupling multiple disciplinary models within a unified and efficient optimisation workflow, **O2**: Develop, adapt, and verify disciplinary models suitable for conceptual design optimisation, achieving an appropriate balance between physical fidelity, validity range, and computational efficiency, in order to enable realistic and scalable FOWT design evaluations. **O3**: Derive and test a novel design metric related to the dynamic behaviour of the FOWT system that simultaneously addresses multiple design requirements, simplifying the optimisation problem and supporting more informed conceptual design decisions.

Together, the research towards these objectives established a coherent framework that advances the efficiency, scope, and realism of FOWT design optimisation at the conceptual stage, leading to the contributions listed in the next section.

8.3 Novel contribution to knowledge

This thesis made three main novel contributions to knowledge, broken down into specific contributions where relevant:

1. Development of a coupled frequency–time domain surrogate-augmented model of floating offshore wind turbine dynamics suitable for assessment of serviceability, ultimate, and fatigue limit states at the conceptual design stage.
 - (a) Modular framework for coupled analysis of FOWT subsystems, including iterative solutions for i) mean platform offset, aero and mooring loads, and ii) viscous damping and platform motion.
 - (b) Evaluation of coupled-solution importance and computational cost, and identification of the most efficient nonlinear solver.
 - (c) Hybrid frequency–time domain model for motions and for ultimate and fatigue loads on tower and moorings.
 - (d) New parametrisation of catenary mooring lines with a shape parameter ensuring compatibility between line length, fairlead, and anchor position.
 - (e) Novel resonance-avoidance constraints based on probability and power content of excitation frequencies, offering greater flexibility than hard natural-frequency limits which enabled quantitative assessment of resonance risk within design optimisation.
 - (f) Formulation of the optimisation problem grounded in offshore standards and suitable for the conceptual design stage.
 - (g) An integrated optimisation of the FOWT support structure, enabling quantitative assessment of performance and cost trade-offs between the design of the tower, platform, and mooring system, with direct comparison of efficiency of the sequential and integrated optimisation approaches.
 - (h) Quantitative characterisation of the trade-off between structural fatigue life (tower and mooring system) and system cost, demonstrating the impact of fatigue constraints on optimal design solutions.

2. Conceptualisation, dataset generation, integration, and verification of a widely applicable time-domain mooring surrogate model, agnostic to platform design and environmental conditions.
3. Investigation of the instantaneous centre of rotation of floating offshore wind turbines, addressing the limitations of common fixed-point assumptions, with development of a numerical model to analyse its statistics, design sensitivities, and design trade-offs.
 - (a) Demonstration that FOWTs do not pitch about a fixed axis, challenging a common misconception.
 - (b) Analytical and time-domain methods to identify ICR statistics and sensitivities to environmental and design parameters.
 - (c) Application of ICR insights to reduce system loads and explore trade-offs in loads across components.

The MDAO framework provided a foundation for efficient integration of disciplinary analyses into a coupled model. It enabled iterative solutions for the mean platform position and loads, showing their significant influence at only a modest computational cost. This, in turn, supported more accurate dynamic response solutions about the converged equilibrium. Integrated within OpenMDAO, the framework separates design variables managed by the optimiser from coupling variables resolved by system solvers. It thus offers a flexible and expandable platform into which custom disciplinary models can be incorporated beyond a single project.

The optimisation problem integrated multiple subsystems of the support structure while incorporating constraints representative of serviceability, ultimate, and fatigue limit states, in a form suitable for the conceptual design stage, when system information is limited. It was made compatible with the efficient, relatively low-fidelity dynamic model developed in this work. The parametrisation scheme ensured compatibility between design variables while preserving design freedom. In addition, a resonance-avoidance constraint acknowledged the stochastic nature of environmental loading, leveraging variability to provide greater flexibility without over-engineering

while maintaining safety. This formulation broadens the feasible design space and improves realism compared with the prevailing practice of imposing strict natural frequency limits.

Motivated by consideration of ICR as an additional, unconventional design metric, the ICR research challenged prevailing misconceptions in the field. Recognising that FOWTs do not pitch about a fixed point enabled the development of a first-principles method to determine the ICR's time history, statistics, and sensitivity to environmental and design parameters. Sensitivity analysis and an application in a later design case study demonstrated load reductions and provided clarity on the method's scope and limitations. The ICR-based metric offers a novel characterisation of FOWT dynamics and highlights opportunities to exploit, rather than resist, the inherent behaviour of floating structures.

The coupled model combines frequency-domain efficiency for motion prediction with time-domain accuracy for structural load assessment, including fatigue evaluation by rainflow counting. Computational bottlenecks were overcome by replacing the BEM and mooring models with surrogates, substantially reducing cost without loss of accuracy. The novel data-driven surrogate for mooring tension uses motion time series at the fairleads and mooring system geometry variables, making it agnostic to platform design and external loads, and thus more widely applicable than existing surrogates. This represents a significant advance over the state of the art, offering a tractable yet realistic means of evaluating mooring performance in conceptual optimisation.

Finally, the optimisation case studies demonstrated the successful integration of the developed framework, models, and optimisation problem formulation within practical design applications. The integrated optimisation of the FOWT support structure showed efficiency gains, enabled broader exploration of the design space, and provided additional insights compared to sequential design approaches, confirming both the feasibility and value of systematic optimisation at the conceptual design stage. The implementation of fatigue constraints in the optimisation highlighted strong interdependencies between the mooring system and platform design, revealing their combined role in governing fatigue response. Lastly, the implementation of the proposed resonance con-

straints within the optimisation workflow demonstrated the ability to navigate a highly non-convex design space efficiently and to make informed trade-offs between competing resonance risks and structural cost, thereby providing effective decision support.

Collectively, these advances establish a pathway to more efficient, integrated, and innovative FOWT design. By overcoming challenges in support structure parametrisation, efficient yet accurate coupled dynamics modelling, and optimisation framework development, this work strengthens confidence in replacing the slow, sequential, and largely manual design process with a systematic, integrated optimisation approach. Such a shift enables broader and faster exploration of the design space, reduces reliance on engineering intuition alone, and allows innovative metrics such as the ICR to inform early-stage decisions. The methods developed not only address immediate challenges in FOWT design but also lay the foundation for more advanced, data-driven, and multidisciplinary approaches in the future. Overall, the thesis demonstrates that systematic, optimisation-based methods can make FOWT conceptual design more scalable, reliable, and adaptable to the challenges of next-generation offshore wind development.

8.4 Future Work

While this thesis advances the FOWT conceptual design optimisation, several directions remain open for further work. These include alternative organisation of the framework, refinements of the optimisation problem formulation and dynamic models, and broader investigations of the ICR to build on the foundations established here.

MDAO framework This thesis employed a Multidisciplinary Feasible (MDF) architecture to manage disciplinary coupling variables separately from the design variables. As one of the simplest and most commonly used formulations, the MDF architecture provides a logical starting point. Future work should investigate alternative formulations, including monolithic (for example, Individual Discipline Feasible, Simultaneous Analysis and Design) and distributed (for example, Collaborative Optimisation, Analytical Target Cascading) architectures [11]. If a more complex design problem is

considered, these may offer improved efficiency, scalability, and flexibility for larger problems. The modular formulation of the coupled model provides a solid foundation for further experimentation with these alternative architectures.

While second-order wave loads were found to be negligible for the OC3-Hywind spar, computing second-order difference- and sum-frequency hydrodynamic loads would be necessary to extend the model's applicability to semisubmersible and TLP floating platforms. These loads were not available from the boundary-element solver used at the beginning of this work, but they can now be obtained from more recent versions of tools such as Nemoh and Capytaine. In general, such second-order loads cannot be incorporated straightforwardly into the present first-order frequency-domain formulation, in which the structure responds at the same frequencies as the incident waves. However, it would be interesting to include the effect of the mean drift load and couple it to the iterative solution for the mean platform position. Because this extended formulation would still rely on first-order hydrodynamics and motion response transfer functions, the overall complexity of the MDA model would increase significantly, making the iterative solution strategies investigated in this thesis particularly valuable.

Optimisation problem formulation While the optimisation problem addressed key metrics related to serviceability, ultimate, and fatigue constraints, a more detailed and closer alignment with requirements from standards and recommended practices could be incorporated to better reflect the real decision-making process. In particular, explicitly representing the design load cases (DLCs) prescribed in the IEC standard would strengthen the link to certification-oriented processes and improve the robustness of conceptual designs. For example, this study selected environmental conditions for analysis in a simplified yet statistically consistent manner; however, more rigorous approaches based on extreme value analysis and fatigue condition binning could be employed in future work to improve the realism of optimisation during the conceptual design stage.

Just as constraints can be formulated more closely to align with standards, the objective functions can also be extended to better reflect the broader goals of floating

wind developers beyond purely economic metrics. As an illustrative example, the total length of the mooring lines could be introduced as a secondary objective within a multi-objective optimisation formulation. This metric would provide a simplified yet meaningful proxy for design aspects related to manufacturability, scalability, and installation efficiency. Although the total chain length was not a primary concern for early-stage wind farms comprising only a limited number of FOWT units, forthcoming large-scale deployments will increasingly necessitate explicit consideration of supply-chain capacity and installation logistics. In this context, the proposed multidisciplinary framework provides a natural pathway to incorporate additional performance metrics, enabling more holistic and industry-relevant design trade-off analyses.

Coupled dynamic model The surrogate model developed for mooring tension and fatigue can be extended to cover a broader range of mooring variables, allowing search beyond the catenary configurations. Verification of the complete coupled model against higher-fidelity models is also essential to distinguish modelling errors from discrepancies due to lower-fidelity physical representations or neglected phenomena. Moreover, the present model is restricted to unidirectional environmental loads with an empirical correction for directionality, a simplification that should be lifted to improve its physical realism and practical applicability.

Instantaneous centre of rotation (ICR) Further work is needed to generalise the conclusions on ICR behaviour beyond the spar and semisubmersible platforms studied here. Future studies should investigate how load characteristics—such as amplitude, point of application, distribution, and frequency—interact with system inertial and restoring properties to affect motions, the ICR, and structural loads. This research should abstract from specific geometries and environmental conditions, focusing instead on the fundamental load–response interactions that govern ICR behaviour. In addition, extending the parametrisation of support structures to include design features such as column separation in semisubmersibles would enable more explicit control of the ICR in design. Evaluating these effects under the design load cases typically considered in feasibility studies would provide practical insight. Finally, developing a more efficient

frequency-domain method to compute the ICR would enable its use in rapid design space exploration and optimisation.

Optimisation study In this work, a local gradient-based optimiser (IPOPT) was employed using a limited number of starting points. While this approach proved effective in identifying improved designs and enabling local exploration of the design space, future work could consider alternative optimisation algorithms and multi-start strategies to improve robustness and increase the likelihood of identifying global optima.

Applicability to different FOWT types This thesis focused on a spar-type FOWT equipped with a 5 MW rotor and a catenary mooring system. While this configuration was industry-relevant at the outset of the study, the field has since shifted toward larger turbines and semisubmersible platforms with more complex mooring arrangements. The design problem formulation and the coupled model should therefore be updated to better reflect current industry practice. For example, employing a single slender-cylinder platform enabled several simplifications, such as neglecting second-order wave loads — an assumption that would not be valid for large-volume semisubmersible platforms. Similarly, the current framework was developed for catenary mooring systems; extending it to include taut or semi-taut configurations, as well as mixed chain-rope layouts, would require reformulating the parameterisation and would substantially increase the complexity of the design space.

Emerging models Finally, the disciplinary models used in this thesis represent only one point in a rapidly evolving modelling ecosystem. As the floating wind research community continues to advance numerical methods and open-source tools, several emerging models could provide more suitable and efficient alternatives for future MDAO studies. These include, but are not limited to:

- Frequency-domain lumped-mass mooring models [162, 285], which would enable direct evaluation of wave-frequency hydrodynamic loads on the mooring lines in the frequency domain. This approach is more accurate and flexible, while also

being efficient, than the approach used here, which struggled with wave-frequency load prediction.

- Second-order frequency-domain hydrodynamic load models [151], which would extend the current framework by allowing accurate representation of low-frequency hydrodynamic loads. These effects are particularly important for large-volume platforms such as semisubmersibles—where second-order forces are significant—while this thesis focused on spars, where they are comparatively less significant.
- Differentiable boundary element method (BEM) models [57, 286], which enable analytical gradient computation and therefore a fully analytical optimisation workflow. The potential-flow solution is often the computational bottleneck in frequency-domain FOWT analysis and numerical differentiation exacerbates this cost. Differentiable BEM models provide gradients analytically and efficiently, offering interpretability benefits beyond those of the surrogate-based approach adopted in this thesis (though at a higher per-evaluation cost than the surrogate).
- Open-source reference controllers [287], which would allow a more realistic representation of turbine–platform control interactions. This thesis relied on prescribing control points from simplified control curves, whereas such controllers can capture full closed-loop dynamics.
- Frequency-domain aerodynamic models [166], which would avoid the need for pre-computation and interpolation of aerodynamic responses from time-domain simulations, as performed here. Integrating these models directly into the frequency-domain formulation would increase transparency and adaptability of the MDAO workflow.

Bibliography

- [1] Mathieu Balesdent, Loïc Brevault, Nathaniel Price, Sebastien Defoort, Rodolphe Le Riche, Nam Kim, Raphael Haftka, and Nicolas Bérend. *Advanced Space Vehicle Design Taking into Account Multidisciplinary Couplings and Mixed Epistemic/Aleatory Uncertainties*. 10 2016.
- [2] Bjorn de Koeijer, Renee Wever, and Jörg Henseler. Realizing Product-Packaging Combinations in Circular Systems: Shaping the Research Agenda. *Packag. Technol. Sci.*, 30(8), 2017.
- [3] J M Jonkman. Dynamics modeling and loads analysis of an offshore floating wind turbine. Technical Report NREL/TP-500-41958, NREL, November 2007.
- [4] T Ashuri, M B Zaaijer, J R R A Martins, G J W van Bussel, and G A M van Kuik. Multidisciplinary design optimization of offshore wind turbines for minimum levelized cost of energy. *Renew. Energy*, 68, 2014.
- [5] Andrew B Lambe and Joaquim R R A Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Struct. Multidiscip. Optim.*, 46(2), 2012.
- [6] Matthew Hall, Brad Buckham, and Curran Crawford. Evolving offshore wind: A genetic algorithm-based support structure optimization framework for floating wind turbines. In *Ocean. 2013 MTS/IEEE Bergen Challenges North. Dimens.*, 2013.

Bibliography

- [7] Christopher Allen, Anthony Viselli, Habib Dagher, Andrew Goupee, Evan Gaertner, Nikhar Abbas, Matthew Hall, and Garrett Barter. Definition of the UMaine VoltturnUS-S Reference Platform Developed for the IEA Wind 15- Megawatt Off-shore Reference Wind Turbine. IEA Wind TCP Task 37. Technical report, 2020.
- [8] Matthew Hall. A Cumulative Multi-Niching Genetic Algorithm for Multimodal Function Optimization. *Int. J. Adv. Res. Artif. Intell.*, 1(9), 2012.
- [9] Gaertner Evan, Jennifer Rinker, Latha Sethuraman, Frederik Zahle, Matt Shields, Roland Feil, Katherine Dykes, Henrik Bredmose, Christopher Allen, Anthony Viselli, Pietro Bortolotti, Benjamin Anderson, Witold Skrzypinski, Garrett Barter, George Scott, Nikhar Abbas, and Fanzhong Meng. IEA Wind TCP Task 37 Definition of the IEA Wind 15-Megawatt Offshore Reference Wind Turbine Technical Report, 2020.
- [10] DNV GL. Dnv-rp-0286:2019 - recommended practice for coupled analysis of floating wind turbines, 2021.
- [11] Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Cambridge University Press, Cambridge, UK, 2021.
- [12] Joyce Lee and Feng Zhao. Global wind report 2021. Technical report, Global Wind Energy Council (GWEC), 2021.
- [13] Anders Myhr, Catho Bjerkseter, Anders Ågotnes, and Tor A. Nygaard. Levelised cost of energy for offshore floating wind turbines in a lifecycle perspective. *Renewable Energy*, 66:714–728, 2014.
- [14] Matilda Kreider, Frank Oteri, Amy Robertson, Chloe Constant, and Elizabeth Gill. Offshore wind energy: Technology below the water. Technical report, National Wind Technology Center, 2022.
- [15] M. Y. Mahfouz, E. Lozon, M. Hall, and P. W. Cheng. Dynamic performance of a passively self-adjusting floating wind farm layout to increase the annual energy production. *Wind Energy Science*, 9(7):1595–1615, 2024.

Bibliography

- [16] Hector del Pozo Gonzalez, Jose Luis Domínguez-García, and Oriol Gomis-Bellmunt. Structural motion control of waked floating offshore wind farms. *Ocean Engineering*, 293:116709, 2024.
- [17] Kaiser Saeed, Jade McMorland, Maurizio Collu, Andrea Coraddu, James Carroll, and David McMillan. Adaptations of offshore wind operation and maintenance models for floating wind. *Journal of Physics: Conference Series*, 2362(1):012036, nov 2022.
- [18] Garrett E Barter, Amy Robertson, and Walter Musial. A systems engineering vision for floating offshore wind cost optimization. *Renew. Energy Focus*, 34(6):1143–1148, 2020.
- [19] Paul Van der Valk and Sven Voormeeren. An overview of modeling approaches for complex offshore wind turbine support structures. 09 2012.
- [20] S Sanchez Perez-Moreno, K Dykes, K O Merz, and M B Zaaier. Multidisciplinary design analysis and optimisation of a reference offshore wind plant. *Journal of Physics: Conference Series*, 1037(4):042004, jun 2018.
- [21] A. Subbulakshmi, Mohit Verma, M. Keerthana, Saptarshi Sasmal, P. Harikrishna, and Santosh Kapuria. Recent advances in experimental and numerical methods for dynamic analysis of floating offshore wind turbines — an integrated review. *Renewable and Sustainable Energy Reviews*, 164:112525, 2022.
- [22] Edris Safavi, Mehdi Tarkian, Johan Ölvander, Hossein Nadali Najafabadi, and Raghu Chaitanya Munjulury. Implementation of collaborative multidisciplinary design optimization for conceptual design of a complex engineering product. *Concurr. Eng. Res. Appl.*, 24(3), 2016.
- [23] Erin E Bachynski and Torgeir Moan. Design considerations for tension leg platform wind turbines. *Mar. Struct.*, 29(1), 2012.
- [24] Jason M Jonkman. Dynamics of offshore floating wind turbines-model development and verification. *Wind Energy*, 12(5), 2009.

Bibliography

- [25] Matthew Hall, Ericka Lozon, Fiona Devoy McAuliffe, Matteo Baudino Bessone, Ilmas Bayati, Malcolm Bowie, Pauline Bozonnet, Michel Castagné, Ju Feng, Stein Housner, Marek Jan Janocha, Zhiyu Jiang, Yong Yook Kim, Donghee Ko, Konstanze Kölle, Chern Fong Lee, Malakonda Reddy Lekkala, Guodong Liang, Mohammad Youssef Mahfouz, Madhan Mohan, Daniel O’Connell, Muk Chen Ong, Julien Prieur, Vishnu R. N. Rajasree, Anja Schnepf, Thor Snedker, Jean Thurston-Keller, and Chris Wright. The iea wind task 49: Reference floating wind array design basis. Technical report, IEA Wind TCP, 2024.
- [26] DNV-RP-C203: Fatigue Design of Offshore Steel Structures. Technical report, DNV, 2024. Recommended Practice.
- [27] S. Butterfield, W. Musial, J. Jonkman, and P. Sclavounos. Engineering challenges for floating offshore wind turbines. Proceedings of the Copenhagen Offshore Wind Conference, 2005.
- [28] Marit I. Kvittem and Torgeir Moan. Time domain analysis procedures for fatigue assessment of a semi-submersible wind turbine. *Marine Structures*, 40:38–59, 2015.
- [29] Yuanchuan Liu, Qing Xiao, Atilla Incecik, Christophe Peyrard, and Decheng Wan. Establishing a fully coupled CFD analysis tool for floating offshore wind turbines. *Renewable Energy*, 112:280–301, 2017.
- [30] Y Huang and D. Wan. Investigation of interference effects between wind turbine and spar-type floating platform under combined wind-wave excitation. *Sustainability*, 12, 2020.
- [31] L.S.P. da Silva, M. de Oliveira, B. Cazzolato, N. Sergiienko, G.A. Amaral, and B. Ding. Statistical linearisation of a nonlinear floating offshore wind turbine under random waves and winds. *Ocean Engineering*, 261, 2022.
- [32] Carlos Eduardo S. Souza and Erin E. Bachynski. Changes in surge and pitch decay periods of floating wind turbines for varying wind speed. *Ocean Engineering*, 180:223–237, 2019.

Bibliography

- [33] P.C.; do Carmo L.H.S.; Alberto I.F.; Malta E.B.; Simos A.N.; Franzini G.R.; Suzuki H.; Gonçalves R.T. Amaral, G.A.; Mello. Seakeeping tests of a fowt in wind and waves: An analysis of dynamic coupling effects and their impact on the predictions of pitch motion response. *Journal of Marine Science and Engineering*, 9, 2021.
- [34] C. Lienard, R. Boisard, and C. Daudin. Aerodynamic behavior of a floating offshore wind turbine. 2019.
- [35] S. Gueydon, I. Bayati, and E.J. de Ridder. Oc6 semisubmersible response to waves in offset positions caused by constant thrust loads. 2021.
- [36] David T. Brown. Chapter 8 - mooring systems. In Subrata K. Chakrabarti, editor, *Handbook of Offshore Engineering*, pages 663–708. Elsevier, London, 2005.
- [37] R. Antonutti, Peyrard C., L. Johanning, A. Incecik, and D. Ingram. The effects of wind-induced inclination on the dynamics of semi-submersible floating wind turbines in the time domain. *Renewable Energy*, 88, 2016.
- [38] L A Schmit. Structural Design by Systematic Synthesis. In *Proceedings, 2nd Conf. Electron. Comput. ASCE*, 1960.
- [39] openmdao.org. OpenMDAO Theory manual. https://openmdao.org/newdocs/versions/latest/theory_manual/theory_manual.html. [accessed 10 September 2021].
- [40] G. Johnson, R. Whittington, D. Angwin, K. Scholes, and P. Regnér. *Exploring Strategy*. Always Learning. Pearson, 2014.
- [41] International Organization for Standardization. En iso 19904-1:2006 petroleum and natural gas industries — floating offshore structures — part 1: Monohulls, semi-submersibles, and spars. ISO Standard, 2006. Section 6.3: Limit states.
- [42] J Jonkman. Definition of the Floating System for Phase IV of OC3. Technical Report NREL/TP-500-47535, NREL, May 2010.

Bibliography

- [43] J Jonkman, S Butterfield, W Musial, and G Scott. Definition of a 5-mw reference wind turbine for offshore system development, 2009.
- [44] Catherine G. Manning. Technology readiness levels, 2023. Accessed: 2024-11-13.
- [45] Paul Sclavounos, Christopher Tracy, and Sungho Lee. Floating offshore wind turbines: Responses in a seastate Pareto optimal designs and economic assessment. In *Proc. Int. Conf. Offshore Mech. Arct. Eng. - OMAE*, volume 6, 2008.
- [46] Ivar Fylling and Petter Andreas Berthelsen. WINDOPT - An optimization tool for floating support structures for deep water wind turbines. In *Proc. Int. Conf. Offshore Mech. Arct. Eng. - OMAE*, volume 5, 2011.
- [47] Anders Myhr and Tor Anders Nygaard. Load reductions and optimizations on tension-leg-buoy offshore wind turbine platforms. In *Proc. Int. Offshore Polar Eng. Conf.*, 2012.
- [48] Meysam Karimi, Matthew Hall, Brad Buckham, and Curran Crawford. A multi-objective design optimization approach for floating offshore wind turbine support structures. *J. Ocean Eng. Mar. Energy*, 3(1), 2017.
- [49] Frank Lemmer, Kolja Müller, Wei Yu, David Schlipf, and Po Wen Cheng. Optimization of floating offshore wind turbine platforms with a self-tuning controller. In *Proc. Int. Conf. Offshore Mech. Arct. Eng. - OMAE*, volume 10, 2017.
- [50] Suguang Dou, Antonio Pegalajar-Jurado, Shaofeng Wang, Henrik Bredmose, and Mathias Stolpe. Optimization of floating wind turbine support structures using frequency-domain analysis and analytical gradients. In *J. Phys. Conf. Ser.*, volume 1618, 2020.
- [51] John Marius Hegseth, Erin E Bachynski, and Joaquim R R A Martins. Integrated design optimization of spar floating wind turbines. *Mar. Struct.*, 72, 2020.
- [52] John Marius Hegseth, Erin E Bachynski, and Joaquim R R A Martins. Design Optimization of Spar Floating Wind Turbines Considering Different Control Strategies. In *J. Phys. Conf. Ser.*, volume 1669, 2020.

Bibliography

- [53] Frank Lemmer, Wei Yu, Kolja Müller, and Po Wen Cheng. Semi-submersible wind turbine hull shape design for a favorable system response behavior. *Mar. Struct.*, 71, 2020.
- [54] John Marius Hegseth, Erin E Bachynski, and Bernt J Leira. Effect of environmental modelling and inspection strategy on the optimal design of floating wind turbines. *Reliability Engineering and System Safety*, 214, 2021.
- [55] Nicol'o Pollini, Antonio Pegalajar-Jurado, Suguang Dou, Henrik Bredmose, and Mathias Stolpe. Gradient-based optimization of a 15 mw wind turbine spar floater. 2021.
- [56] P Rohrer, E E Bachynski-Polić, and M Collette. Towards gradient-based design optimization of fully-flexible tension-leg platform wind turbines. *Journal of Physics: Conference Series*, 2362(1), nov 2022.
- [57] *Gradient-Based Design Optimization of Fully-Flexible Floating Wind Turbines Using Modal Analysis*, volume Volume 8: Ocean Renewable Energy of *International Conference on Offshore Mechanics and Arctic Engineering*, 2023.
- [58] Nicolò Pollini, Antonio Pegalajar-Jurado, and Henrik Bredmose. Design optimization of a TetraSpar-type floater and tower for the IEA Wind 15 MW reference wind turbine. *Marine Structures*, 90:103437, 2023.
- [59] Nikhar J. Abbas, John Jasa, Daniel S. Zalkind, Alan Wright, and Lucy Pao. Control co-design of a floating offshore wind turbine. *Applied Energy*, 353, 2024.
- [60] R G Standing, W J Brendling, and D Wilson. Recent developments in the analysis of wave drift forces, low-frequency damping and response. In *Proc. Annu. Offshore Technol. Conf.*, volume 1987-April, 1987.
- [61] H Lie and N Sødahl. Simplified Dynamic Model for Estimation of Extreme Anchorline Tension, 1993.
- [62] M K Al-Solihat and M Nahon. Stiffness of slack and taut moorings. *Ships and Offshore Structures*, 11:890–904, 2016.

Bibliography

- [63] Joaquim R. R. A. Martins and Andrew B. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075, 2013.
- [64] Steven R Hirshorn. *NASA Systems Engineering Handbook*. 2 edition, 2007.
- [65] James Dobbin, David Quarton, Joe Phillips, and Paul Reynolds. Project force: Offshore wind cost reduction through integrated design, 2014.
- [66] S S Perez-Moreno, M B Zaaier, C L Bottasso, K Dykes, K O Merz, P-E Rethore, and F Zahle. Roadmap to the multidisciplinary design analysis and optimisation of wind energy systems. 2016.
- [67] C. L. Bottasso, F. Campagnolo, and A. Croce. Multi-disciplinary constrained optimization of wind turbines. *Multibody System Dynamics*, 27:21–53, 2012.
- [68] Pietro Bortolotti, Carlo L. Bottasso, and Alessandro Croce. Combined preliminary–detailed design of wind turbines. *Wind Energ. Sci.*, 1:71–88, 2016.
- [69] B Grossman, Z Gurdal, R T Haftka, G J Strauch, and W M Eppard. Integrated aerodynamic/structural design of a sailplane wing. *J. Aircr.*, 25(9), 1988.
- [70] G F Clauss and L Birk. Hydrodynamic shape optimization of large offshore structures. *Appl. Ocean Res.*, 18(4), 1996.
- [71] Beom Seon Jang, Jae Dong Kim, Tae Yoon Park, and Sang Bae Jeon. FEA based optimization of semi-submersible floater considering buckling and yield strength. *Int. J. Nav. Archit. Ocean Eng.*, 11(1), 2019.
- [72] Toshio Sugita and Hideyuki Suzuki. A study on TLP hull sizing by utilizing optimization algorithm. *J. Mar. Sci. Technol.*, 21(4), 2016.
- [73] Christopher Tracy. Parametric Design of Floating Wind Turbines. Master’s thesis, Massachusetts Institute of Technology, 2007.
- [74] Ankur Sinha and Kalyanmoy Deb. Bilevel multi-objective optimization and decision making. *Stud. Comput. Intell.*, 482, 2013.

Bibliography

- [75] Onur Bilgen, Joaquim R.R.A. Martins, Andrew Ning, Laurent Burlion, Andy Platt, and et al. Openturbinecode (otcd): An open-source floating offshore wind turbine multidisciplinary control co-design optimization framework. In *NAWEA/WindTech 2024*, 2024. Conference paper.
- [76] Xianping Du, Jinbin Liang, Guowei Qian, Yang Yang, Peng Xie, and Kai Zhang. A computational framework for the co-optimization of platform hydrodynamic and passive structural control of floating offshore wind turbines. *Ocean Engineering*, 293, 2024.
- [77] W. Yu, S. T. Zhou, F. Lemmer, and P. W. Cheng. Control co-design optimization of floating offshore wind turbines with tuned liquid multi-column dampers. *Wind Energy Science*, 9(4), 2024.
- [78] Pier Davide Ciampa and Björn Nagel. Towards the 3rd generation MDO collaboration environment. In *30th Congr. Int. Council. Aeronaut. Sci. ICAS*, 2016.
- [79] Justin S. Gray, John T. Hwang, Joaquim R.R.A. Martins, Kenneth T. Moore, and Bret A. Naylor. Openmdao: an open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59, 2019.
- [80] Michael S Eldred, Sand Report, Anthony A Giunta, Bart G Van Bloemen Waanders, Steven F Wojtkiewicz, William E Hart, and Mario P Alleva. DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis (SAND 2001-3514). *Analysis*, (April), 2002.
- [81] Ping He, Charles A Mader, Joaquim R R A Martins, and Kevin J Maki. DA Foam: An open-source adjoint framework for multidisciplinary design optimization with openFOAM. *AIAA J.*, 58(3), 2020.
- [82] openmdao.org. OpenMDAO 1.7.3 Alpha documentation. <https://openmdao.readthedocs.io/en/1.7.3/tags/Multiprocessing.html>. [accessed 15 March 2021].

Bibliography

- [83] Ping He, Grzegorz Filip, Joaquim R R A Martins, and Kevin J Maki. Design optimization for self-propulsion of a bulk carrier hull using a discrete adjoint method. *Computers and Fluids*, 192, 2019.
- [84] C C Xia, Y J Gou, S H Li, W F Chen, and Chongzhong Shao. An automatic aerodynamic shape optimisation framework based on dakota. 2018.
- [85] Jeremy Agte, Olivier De Weck, Jaroslaw Sobieszczanski-Sobieski, Paul Arendsen, Alan Morris, and Martin Spieck. MDO: Assessment and direction for advancement-an opinion of one international group. *Struct. Multidiscip. Optim.*, 40(1-6), 2010.
- [86] Jamshid A Samareh. A survey of shape parametrization techniques. In *CEAS/AIAA/ICASE/NASA Langley Int. Forum Aeroelasticity Struct. Dyn.*, pages 333–343. NASA Langley Research Center, 1999.
- [87] Jamshid A Samareh. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *American Institute of Aeronautics and Astronautics Journal*, 39, 2001.
- [88] Matthias Brommundt, Ludwig Krause, Karl Merz, and Michael Muskulus. Mooring system optimization for floating wind turbines using frequency domain analysis. In *Energy Procedia*, volume 24, 2012.
- [89] Ajit C Pillai, Philipp R Thies, and Lars Johanning. Mooring system design optimization using a surrogate assisted multi-objective genetic algorithm. *Eng. Optim.*, 51(8), 2019.
- [90] Matthew Hall, Brad Buckham, and Curran Crawford. Hydrodynamics-based floating wind turbine support platform optimization: A basis function approach. *Renew. Energy*, 66, 2014.
- [91] M Leimeister, A Kolios, and M Collu. Critical review of floating support structures for offshore wind farm deployment. In *J. Phys. Conf. Ser.*, volume 1104, 2018.

Bibliography

- [92] Adebayo Ojo, Maurizio Collu, and Andrea Coraddu. Parametrisation scheme for multidisciplinary design analysis and optimisation of a floating offshore wind turbine substructure – oc3 5mw case study. *Journal of Physics: Conference Series*, 2265:042009, 05 2022.
- [93] Jean Christophe Gilloteaux and Pauline Bozonnet. Parametric analysis of a cylinder-like shape floating platform dedicated to multi-megawatt wind turbine. In *Proc. Int. Offshore Polar Eng. Conf.*, 2014.
- [94] T Ashuri, J R R A Martins, M B Zaaier, G A M van Kuik, and G J W van Bussel. Aeroservoelastic design definition of a 20 MW common research wind turbine model. *Wind Energy*, 19(11), 2016.
- [95] L Castro-Santos, M deCastro, X Costoya, A Filgueira-Vizoso, I Lamas-Galdo, A Ribeiro, J.M Dias, and M Gómez-Gesteira. Economic feasibility of floating offshore wind farms considering near future wind resources: Case study of iberian coast and bay of biscay. *Int. J. Environ. Res. Public Health*, 18, 2021.
- [96] T. Stehly, P. Beiter, and P. Duffy. 2019 cost of wind energy review. Tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States), 2020.
- [97] Rhodri James and Marc Costa Ros. Floating Offshore Wind : Market and Technology Review. Technical report, Carbon Trust, June 2015.
- [98] E E Bachynski. Fixed and floating offshore wind turbine support structures, 2018.
- [99] Amir R Nejad, Erin E Bachynski, and Torgeir Moan. Effect of axial acceleration on drivetrain responses in a spar-type floating wind turbine. *J. Offshore Mech. Arct. Eng.*, 141(3), 2019.
- [100] L Pustina, C Lugni, G Bernardini, J Serafini, and M Gennaretti. Control of power generated by a floating offshore wind turbine perturbed by sea waves. *Renew. Sustain. Energy Rev.*, 132, 2020.

Bibliography

- [101] Olimpo Anaya-Lara, John Olav Tande, Kjetil Uhlen, and Karl Merz. *Offshore Wind Energy Technology*. 2018.
- [102] H Namik and K Stol. Individual blade pitch control of floating offshore wind turbines. *Wind Energy*, 13(1), 2010.
- [103] Jun B Rho, Hang S Choi, Hyun S Shin, and In K Park. A study on Mathieu-type instability of conventional spar platform in regular waves. *Int. J. Offshore Polar Eng.*, 15(2), 2005.
- [104] Herbj0rn Alf Haslum and Odd M Faltinsen. Simplified methods applied to non-linear motion of spar platforms (A thesis submitted in partial fulfilment of the requirements for the degree of Doktor Ingeniør. Trondheim), 2000.
- [105] Kai-Tung Ma, Yong Luo, Thomas Kwan, and Yongyan Wu. Chapter 6 - fatigue analysis. In *Mooring System Engineering for Offshore Structures*, pages 115–137. Gulf Professional Publishing, 2019.
- [106] Jijian Lian, Huan Zhou, and Xiaofeng Dong. A Theoretical Approach for Resonance Analysis of Wind Turbines under 1P/3P Loads. *Energies*, 15, 2022.
- [107] Vincenzo Piscopo, Antonio Scamardella, Giovanni Battista Rossi, Francesco Crenna, and Marta Berardengo. Fatigue assessment of moorings for floating offshore wind turbines by advanced spectral analysis methods. *Journal of Marine Science and Engineering*, 10, 2022.
- [108] Det Norske Veritas (DNV). Design of Floating Wind Turbine Structures, 2018.
- [109] Innosea, Cobra, Esteyco, DTU, JDR, UPC, FIHAC, UL INT GMBH, and US-TUTTM. Corewind project - wp1: D1.2 design basis. Corewind Project Report, 2019. 2019-2023.
- [110] L Meriam, J and G Kraige, L. *Engineering Mechanics, Volume 2, Dynamics*. John Wiley & Sons, Inc., 7th edition, 1993.

Bibliography

- [111] Antonio Carlos Fernandes, Peyman Asgari, and Anderson R.W. Soares. Asymmetric roll center of symmetric body in beam waves. *Ocean Engineering*, 112:66–75, 2016.
- [112] Paulo Esperança, Joel Sales, Stergios Liapis, João Paulo Matsuura, and Wes Schott. An Experimental Investigation of Roll Motions of an FPSO. volume 1, 2008.
- [113] Binrong Wen, Zhan-Wei Li, Zhihao Jiang, Xinliang Tian, Xing-Jian Dong, and Z. Peng. Floating wind turbine power performance incorporating equivalent turbulence intensity induced by floater oscillations. *Wind Energy*, 25, 2022.
- [114] Thanhtoan Tran, Donghyun Kim, and Jinseop Song. Computational Fluid Dynamic Analysis of a Floating Offshore Wind Turbine Experiencing Platform Pitching Motion. *Energies*, 7:5011–5026, 2014.
- [115] Hannah M Johlas, Luis A Martínez-Tossas, Matthew J Churchfield, Matthew A Lackner, and David P Schmidt. Floating platform effects on power generation in spar and semisubmersible wind turbines. *Wind Energy*, 2021.
- [116] Felix Kelberlau. Free decay testing of a semisubmersible offshore floating platform for wind turbines in model scale (Master’s thesis).
- [117] W. P. Stewart and W. A. Ewers. *Wave Induced Motions of Marine Deck Cargo Barges*. Imperial College Press, London, 1979. Second International Conference on Behaviour of Off-Shore Structures, BOSS’79.
- [118] R.G. Standing. Prediction Of Viscous Roll Damping And Response Of Transportation Barges In Waves. International Ocean and Polar Engineering Conference, 1991.
- [119] F.S. Souza, R.S. Anderson, A.C. Oliveira, and A.C. Fernandes. Influência do Centro de Rotação através de Ensaios de Decaimento no Movimento de Balanço Transversal de Cascos Típicos de FPSO (in Portuguese). XXIV Congresso Nacional de Transportes Marítimos e Construção Naval e Offsh, 2012.

Bibliography

- [120] Daniel de Oliveira Costa, Antonio Carlos Fernandes, Joel Sena Sales Junior, and Peyman Asgari. Instantaneous Center of Rotation in Pitch Response of a FPSO Submitted to Head Waves. volume Volume 7A: Ocean Engineering of *International Conference on Offshore Mechanics and Arctic Engineering*, 2018.
- [121] Daniel de Oliveira Costa, Antonio Carlos Fernandes, and Joel Sena Sales Junior. Further study on the instantaneous rotation center in pitch and its distribution in space for a moored vessel submitted to head regular waves. *Ocean Engineering*, 218, 2020.
- [122] Daniel de Oliveira Costa, Joel Sena Sales Junior, and Antonio Carlos Fernandes. Instantaneous Center of Rotation of a Vessel Submitted to Oblique Waves. volume Volume 6B: Ocean Engineering of *International Conference on Offshore Mechanics and Arctic Engineering*, 2020.
- [123] Lene Eliassen. Aerodynamic loads on a wind turbine rotor in axial motion. <https://core.ac.uk/display/30846862?source=4>.
- [124] Mert Kaptan, Bjørn Skaare, Zhiyu Jiang, and Muk Chen Ong. Analysis of spar and semi-submersible floating wind concepts with respect to human exposure to motion during maintenance operations. *Marine Structures*, 83, 2022.
- [125] Carlos Eduardo Silva de Souza and Erin E. Bachynski-Polić. Design, structural modeling, control, and performance of 20 mw spar floating wind turbines. *Marine Structures*, 84:103182, 2022.
- [126] Victor Benifla and Frank Adam. Development of a genetic algorithm code for the design of cylindrical buoyancy bodies for floating offshore wind turbine substructures. *Energies*, 15(3), 2022.
- [127] Michael J D Powell. The bobyqa algorithm for bound constrained optimization without derivatives, 2009.

Bibliography

- [128] USTUTT. Qualification of innovative floating substructures for 10MW wind turbines and water depths greater than 50m. Optimisation framework and methodology for optimized floater design, 2016.
- [129] Kapil Khanal, Carlos A. Michelén Ströfer, Matthieu Ancellin, and Maha N. Haji. Fully differentiable boundary element solver for hydrodynamic sensitivity analysis of wave-structure interactions. *Applied Ocean Research*, 163:104707, 2025.
- [130] Peter J. Rohrer and Erin E. Bachynski-Polić. Analytical gradients of first-order diffraction and radiation forces for design optimization of floating structures. *Applied Ocean Research*, 152:104198, 2024.
- [131] Joaquim R R A Martins and John T Hwang. Review and unification of methods for computing derivatives of multidisciplinary computational models. In *AIAA J.*, volume 51, 2013.
- [132] Philip E Gill, Walter Murray, and Michael A Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization, 2005.
- [133] National Renewable Energy Laboratory (NREL). OpenFAST Documentation. Release 1.0. Technical report, 2019.
- [134] Andrew Cordle and Jason Jonkman. State of the art in floating wind turbine design tools. In *Proc. Int. Offshore Polar Eng. Conf.*, 2011.
- [135] Amy N Robertson, Fabian Wendt, Jason M Jonkman, Wojciech Popko, Habib Dagher, Sebastien Gueydon, Jacob Qvist, Felipe Vittori, José Azcona, Emre Uzunoglu, Carlos Guedes Soares, Rob Harries, Anders Yde, Christos Galinos, Koen Hermans, Jacobus Bernardus De Vaal, Pauline Bozonnet, Ludovic Bouy, Ilmas Bayati, Roger Bergua, Josean Galvan, Iñigo Mendikoa, Carlos Barrera Sanchez, Hyunkyoung Shin, Sho Oh, Climent Molins, and Yannick Debruyne. OC5 Project Phase II: Validation of Global Loads of the DeepCwind Floating Semisubmersible Wind Turbine. In *Energy Procedia*, volume 137, 2017.

Bibliography

- [136] Lu Wang, Amy Robertson, Jason Jonkman, Yi-Hsiang Yu, Arjen Koop, Adrià Borràs Nadal, Haoran Li, Erin Bachynski-Polić, Romain Pinguete, Wei Shi, Xinmeng Zeng, Yang Zhou, Qing Xiao, Rupesh Kumar, Hamid Sarlak, Edward Ransley, Scott Brown, Martyn Hann, Stefan Netzband, Malwin Wermbter, and Beatriz Méndez López. Oc6 phase ib: Validation of the cfd predictions of difference-frequency wave excitation on a fowt semisubmersible. *Ocean Engineering*, 241, 2021.
- [137] J M Jonkman and M L Buhl. Loads Analysis of a Floating Offshore Wind Turbine Using Fully Coupled Simulation. Technical Report NREL/CP-500-41714, NREL, 2007.
- [138] Richard Lupton. Frequency-domain modelling of floating wind turbines, 2014.
- [139] Madjid Karimirad. Stochastic dynamic response analysis of spar-type wind turbines with catenary or taut mooring systems, 2011.
- [140] M. Folley. Chapter 4 - spectral-domain models. In Matt Folley, editor, *Numerical Modelling of Wave Energy Converters*, pages 67–80. Academic Press, 2016.
- [141] G. Ferri, E. Marino, N. Bruschi, and C. Borri. Platform and mooring system optimization of a 10 MW semisubmersible offshore wind turbine. *Renewable Energy*, 182, 2022.
- [142] NREL. Wind Energy with Integrated Servo-controls Toolset (WEIS), 2024.
- [143] NREL. ROSCO Version 2.0.0, 2020.
- [144] Minoo H Patel. *Dynamics of Offshore Structures*. Butterworths, 1989.
- [145] Peter Joshua Rohrer and Erin Elizabeth Bachynski-Polic. Analytical gradients of first-order diffraction and radiation forces for design optimization of floating structures. *Applied Ocean Research*, 152, 2024.
- [146] F Rajabi, S Ghosh, C Oran, and D R Shield. Review of Semisubmersible and Tension Leg Platform Analysis Techniques, Volume 1 - Literature Survey. TM No.: 44-85-02CR, 1985.

Bibliography

- [147] Tiago. Duarte, A. Sarmento, and Jason. Jonkman. Effects of second-order hydrodynamic forces on floating offshore wind turbines. In *AIAA SciTech Forum*, number AIAA 2014-0361, 2014.
- [148] I Bayati, J Jonkman, A Robertson, and A Platt. The effects of second-order hydrodynamics on a semisubmersible floating offshore wind turbine. 2014.
- [149] J. Jonkman. Definition of the Floating System for Phase IV of OC3. Technical report, National Renewable Energy Laboratory (NREL), 2010.
- [150] Line Roald, Jason Jonkman, Amy Robertson, and Ndaona Chokani. The effect of second-order hydrodynamics on floating offshore wind turbines. *Energy Procedia*, 35, 2013. DeepWind’2013 – Selected papers from 10th Deep Sea Offshore Wind RD Conference, Trondheim, Norway, 24 – 25 January 2013.
- [151] Ruddy Kurnia and Guillaume Ducrozet. NEMOH: Open-source boundary element solver for computation of first- and second-order hydrodynamic loads in the frequency domain. *Computer Physics Communications*, 292:108885, 2023.
- [152] M Karch, M Voßbeck, F Wendt, L M Guindo, F Borisade, L Wilms, A B Aardal, N Psychogios, G Grimsrud, and D Matha. Approaches and challenges in feed and detailed design process of floating substructures. *Journal of Physics: Conference Series*, 2875(1), 2024.
- [153] Frank Sandner, David Schlipf, Denis Matha, Robert Seifried, and Po Wen Cheng. Reduced Nonlinear Model of a Spar-mounted Floating Wind Turbine. 2012.
- [154] David Schlipf, Dominik Johannes Schlipf, and Martin Kühn. Nonlinear model predictive control of wind turbines using LIDAR. In *Wind Energy*, volume 16, 2013.
- [155] Alexander J Coulling, Andrew J Goupee, Amy N Robertson, and Jason M Jonkman. Importance of second-Order difference-Frequency wave-Diffraction forces in the validation of a fast semi-Submersible floating wind turbine model. In *Proc. Int. Conf. Offshore Mech. Arct. Eng. - OMAE*, volume 8, 2013.

Bibliography

- [156] Erin E Bachynski and Torgeir Moan. Ringing loads on tension leg platform wind turbines. *Ocean Eng.*, 84, 2014.
- [157] J N Newman. Second-order, slowly-varying forces on vessels in irregular waves. Technical Report MIT Paper 19, Massachusetts Institute of Technology, 1975.
- [158] Frank Lemmer. Low-order modeling, controller design and optimization of floating oshore wind turbines, 2018.
- [159] Leon E. Borgman. The spectral density for ocean wave forces, 1967.
- [160] Det Norske Veritas. Environmental Conditions And Environmental Loads, Recommended Practice DNV-RP-C205, 2010.
- [161] Kjell Larsen and Peter Sandvik. Efficient methods for the calculation of dynamic mooring line tension. In *Proceedings of the First ISOPE European Offshore Mechanics Symposium (EUROMS-90)*, Trondheim, Norway, 1990. Paper ISOPE-E-90-058.
- [162] Serag-Eldin Abdelmoteleb and Erin E. Bachynski-Polić. A frequency-domain optimization procedure for catenary and semi-taut mooring systems of floating wind turbines. *Marine Structures*, 101:103768, 2025.
- [163] Garrad Hassan. DNV GL White Paper on Definitions of Availability Terms for the Wind Industry. Document No.: EAA-WP-15, 2017.
- [164] Tony Burton, Nick Jenkins, David Sharpe, and Ervin Bossanyi. *Wind Energy Handbook, Second Edition*. 2011.
- [165] Frank Sandner, David Schlipf, Denis Matha, and Po Wen Cheng. Integrated optimization of floating wind turbine systems. In *Proc. Int. Conf. Offshore Mech. Arct. Eng. - OMAE*, volume 9B, 2014.
- [166] Serag-Eldin Abdelmoteleb and Erin E. Bachynski-Polić. Evaluating different approaches for modelling rotor aero-servo-dynamics in frequency-domain analysis of floating wind turbines. *Wind Energy*, 28(11):e70060, 2025.

Bibliography

- [167] A Halfpenny. Dynamic analysis of both on and offshore wind turbines in the frequency domain. Technical report, 1998.
- [168] Matthew Hall, Ericka Lozon, Stein Housner, and Senu Srinivas. Design and analysis of a ten-turbine floating wind farm with shared mooring lines. *Journal of Physics: Conference Series*, 2362(1), 2022.
- [169] J. Jasa, P. Bortolotti, D. Zalkind, and G. Barter. Effectively using multifidelity optimization for wind turbine design. *Wind Energy Science*, 7(3), 2022.
- [170] Alexander I. J. Forrester, András Sóbester, and Andy J. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide*. John Wiley & Sons, Chichester, 2008.
- [171] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- [172] Nestor V. Queipo, Raphael T. Haftka, Wei Shyy, Tushar Goel, Rajkumar Vaidyanathan, and Kevin Tucker. Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41:1–28, 2005.
- [173] Xinshu Zhang, Xingyu Song, Wenzhen Qiu, Zhiming Yuan, Yunxiang You, and Naiming Deng. Multi-objective optimization of tension leg platform using evolutionary algorithm based on surrogate model. *Ocean Engineering*, 148, 2018.
- [174] Wenzhen Qiu, Xingyu Song, Kaiyuan Shi, Xinshu Zhang, Zhiming Yuan, and Yunxiang You. Multi-objective optimization of semi-submersible platforms using particle swarm optimization algorithm based on surrogate model. *Ocean Engineering*, 178, 2019.
- [175] Chun Bao Li and Joonmo Choung. Fatigue damage analysis for a floating offshore wind turbine mooring line using the artificial neural network approach. *Ships and Offshore Structures*, 12(sup1):S288–S295, 2017.

Bibliography

- [176] H. Zhang, H. Wang, X. Cai, J. Xie, Y. Wang, and N. Zhang. Novel method for designing and optimising the floating platforms of offshore wind turbines. *Ocean Engineering*, 266, 2022.
- [177] Pau Trubat, Adrián Herrera, and Climent Molins. Mooring optimization using ml techniques. In *Proceedings of the ASME 2022 International Offshore Wind Technical Conference (IOWTC 2022)*. American Society of Mechanical Engineers (ASME), 2022.
- [178] X. Ye, P. Zheng, D. Qiao, X. Zhao, Y. Zhou, and L. Wang. Multi-objective optimization design of a mooring system based on the surrogate model. *Journal of Marine Science and Engineering*, 12(10), 2024.
- [179] Hongwei Wang, Qingao Ran, Gang Ma, Jin Wen, Jianhua Zhang, and Sulian Zhou. Optimization design of floating offshore wind turbine mooring system based on dnn and nsga-iii. *Ocean Engineering*, 316:119915, 2025.
- [180] M. Baudino Bessone, D. Singh, T. Kalimeris, E. Bachynski-Polić, and A. Viré. Surrogate-assisted optimization of floating wind turbine substructure. *Journal of Physics: Conference Series*, 2767:062032, 2024.
- [181] NREL. pyhams: A python-based hybrid analysis model for simulation of offshore wind turbines, 2024. Accessed: 2024-12-02.
- [182] Matthew Hall, Stein Housner, Daniel Zalkind, Pietro Bortolotti, David Ogden, and Garrett Barter. An open-source frequency-domain model for floating wind turbine design optimization. *Journal of Physics: Conference Series*, 2265:042020, 2022.
- [183] Peter J. Rohrer, Erin E. Bachynski-Polić, and John Marius Hegseth. Gradient-based design optimization of semi-submersible floating wind turbines. *Marine Structures*, 106:103951, 2026.
- [184] Andrea Coraddu, Luca Oneto, Miltos Kalikatzarakis, Davide Ilardi, and Maurizio Collu. Floating Spar-Type Offshore Wind Turbine Hydrodynamic Response

Bibliography

- Characterisation: a Computational Cost Aware Approach. In *Glob. Ocean. 2020 Singapore – U.S. Gulf Coast*, pages 1–8, 2020.
- [185] Theodoros Kalimeris. Floating wind turbine optimization via surrogate modelling. Master’s thesis, Norwegian University of Science and Technology, 2023.
- [186] R Lano. The N2 Chart. *TRW Software Series*, 1977.
- [187] Xiaoming Ran, Vincent Leroy, and Erin E. Bachynski-Polić. Hydroelastic response of a flexible spar floating wind turbine: Numerical modelling and validation. *Ocean Engineering*, 286:115635, 2023.
- [188] P.J. Moriarty and A.C. Hansen. AeroDyn Theory Manual. NREL/TP-500-36881, 2005.
- [189] International Electrotechnical Commission (IEC). IEC 61400-1 International Standard. Wind energy generation systems – Part 1: Design requirements, 2019.
- [190] Matthew Hall, Scott Housner, Senu Srinivas, and Stuart Wilson. MoorPy: Quasi-Static Mooring Analysis in Python. Technical report, National Renewable Energy Laboratory, Golden, CO, 2021.
- [191] J. N. Newman. *Marine Hydrodynamics, 40th Anniversary Edition*. The MIT Press Cambridge, Massachusetts, 2 edition, 2018.
- [192] Jack A Myers. Handbook of equations for mass and area properties of various geometrical shapes, 1962.
- [193] Matthieu Ancellin and Frédéric Dias. Capytaine: a python-based linear potential flow solver. *The Journal of Open Source Software*, 4, 2019.
- [194] C.E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
- [195] Erin Elizabeth Bachynski. *Design and Dynamic Analysis of Tension Leg Platform Wind Turbines*. PhD thesis, Norwegian University of Science and Technology, Trondheim, 2014.

Bibliography

- [196] K Patryniak and Collu. Multidisciplinary design analysis and optimisation of floating offshore wind turbine support structures - coupled model solution strategies. *Journal of Physics Conference Series*, 2626, 2023.
- [197] C.G.J. Jacobi. Über ein leichtes verfahren die in der theorie der säcularstörungen vorkommenden gleichungen numerisch aufzulösen. *Theorie der Sucularstorungen*, 1846. , 1846.
- [198] L Seidel. On a process for solving by successive approximation the equations to which the method of least squares leads as well as linear equations generally. *Abhandlungen der Mathematisch-Physikalischen Klasse der Königlich Bayerischen Akademie der Wissenschaften*, 11:81–108, 1874.
- [199] B M Irons and R C Tuck. A version of the Aitken accelerator for computer iteration. *International Journal for Numerical Methods in Engineering*, 1:275–277, 1969.
- [200] C. G. Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of Computation*, 19:469–531, 1965.
- [201] DNV-ST-0119:2021 Floating Wind Turbine Structures. Standard, DNV – Det Norske Veritas, 2021. Floating Wind Turbine Structures.
- [202] IEC 61400-3-2: Wind energy generation systems – Part 3-2: Design requirements for floating offshore wind turbines, 2025.
- [203] DNV-RP-C205: Environmental Conditions and Environmental Loads. Technical report, 2021.
- [204] LR-RP-003-2024: Recommended Practice for Floating Offshore Wind Turbine Support Structures, 2024.
- [205] Mohammad Youssef Mahfou, Mohammad Salari, Sergio Hernández, Fernando Vigara, Climent Molins, Pau Trubat, Henrik Bredmose, and Antonio Pegalajar-Jurado. D1.3. Public design and FAST models of the two 15MW floater-turbine concepts. Technical report, CoreWind Project, 2020.

Bibliography

- [206] Frederik Zahle, Athanasios Barlas, Kenneth Lønbæk, Pietro Bortolotti, Daniel Zalkind, Lu Wang, Casper Labuschagne, Latha Sethuraman, and Garrett Barter. *Definition of the IEA Wind 22-Megawatt Offshore Reference Wind Turbine*. Technical University of Denmark, 2024. DTU Wind Energy Report E-0243 IEA Wind TCP Task 55.
- [207] Vicinay Marine. Vicinay Marine — Mooring chains and floating mooring solutions. <https://www.vicinaymarine.com/>. Accessed: 2025-09-10.
- [208] NREL. Moorprops.py, 2025. Accessed: 2025-04-09.
- [209] Scottish Government. Sectoral Marine Plan for Offshore Wind Energy. Technical report, 2020.
- [210] IEC. Technical Specification IEC 61400-3:2009. Wind turbines - Part 3: Design requirements for offshore wind turbines, 2009.
- [211] T. M. Fulton and W. P. Stewart. Vertical loads on drag embedment anchors. In *Proceedings of the 26th Annual Offshore Technology Conference (OTC 7491)*, Houston, Texas, USA, May 1994. Offshore Technology Conference.
- [212] DNVGL-RP-C202: Buckling Strength of Shells. Technical report, DNV GL, 2002. Recommended Practice.
- [213] Det Norske Veritas (DNV). Offshore standards – dnv-os-e301. position mooring, 2021.
- [214] IEC. Technical Specification IEC 61400-3:2009. Wind turbines – Part 3: Design requirements for offshore wind turbines. Technical report, International Electrotechnical Commission, 2009.
- [215] J.H. Vugts. Fatigue damage assessments and the influence of wave directionality. *Applied Ocean Research*, 27(3):173–185, 2005.
- [216] Offshore Standards — DNV-OS-C101: Structural design of offshore units, 2024. Incorporates Amendment July 2025.

Bibliography

- [217] K. Patryniak, M. Collu, J. Jonkman, M. Hall, G. Barter, D. Zalkind, and A. Coraddu. Investigation into the instantaneous centre of rotation for enhanced design of floating offshore wind turbines. *Wind Energy Science*, 10(9):2051–2077, 2025.
- [218] NREL. Openfast, v3.5. <https://github.com/openfast/openfast/>, 2024.
- [219] J. M. Jonkman. The new modularization framework for the fast wind turbine cae tool. In *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, pages 7–10, Grapevine (Dallas/Ft. Worth Region), TX, January 2013. American Institute of Aeronautics and Astronautics. AIAA-2013-0202, NREL/CP-5000-57228, Golden, CO: National Renewable Energy Laboratory.
- [220] C. H. Lee and J. N. Newman. WAMIT User Manual, Versions 6.4. Technical report, Massachusetts Institute of Technology, 2006.
- [221] OpenFAST developers. rtest: Regression test data for OpenFAST. <https://github.com/OpenFAST/r-test>, 2025. Accessed: 2025-08-16.
- [222] Marco D. Masciola. MAP++ Mooring Analysis Program. <https://www.nrel.gov/wind/nwtc/map-plus-plus.html>. [accessed 13/12/2022].
- [223] Marco Masciola, Jason Jonkman, and Amy Robertson. Implementation of a Multisegmented, Quasi-Static Cable Model. International Ocean and Polar Engineering Conference, 2013.
- [224] Matthew Hall and Andrew Goupee. Validation of a lumped-mass mooring line model with DeepCwind semisubmersible model test data. *Ocean Engineering*, 104, 2015.
- [225] IEC. Technical Specification IEC 61400-3-2. Wind energy generation systems - Part 3-2: Design requirements for floating offshore wind turbines, 2019.
- [226] J. Jonkman. Definition of the Floating System for Phase IV of OC3. Technical report, National Renewable Energy Laboratory (NREL), 2010.

Bibliography

- [227] K. Stewart. Expected Value. Encyclopedia Britannica. <https://www.britannica.com/topic/expected-value>. [accessed 24 October 2022].
- [228] J. L. Jr. Hodges. The significance probability of the smirnov two-sample test. *Arkiv för Matematik*, 3(43):469–486, 1958.
- [229] A N Robertson, S Gueydon, E Bachynski, L Wang, J Jonkman, D Alarcón, E Amet, A Beardsell, P Bonnet, B Boudet, C Brun, Z Chen, M Féron, D Forbush, C Galinos, J Galvan, P Gilbert, J Gómez, V Harnois, F Haudin, Z Hu, J Le Dreff, M Leimeister, F Lemmer, H Li, G Mckinnon, I Mendikoa, A Moghtadaei, S Netzbant, S Oh, A Pegalajar-Jurado, M Q Nguyen, K Ruehl, P Schünemann, W Shi, H Shin, Y Si, F Surmont, P Trubat, J Qwist, and S Wohlfahrt-Laymann. OC6 Phase I: Investigating the underprediction of low-frequency hydrodynamic loads and responses of a floating wind turbine. *Journal of Physics: Conference Series*, 1618(3), 2020.
- [230] Arjun Srinivas, Bryson Robertson, Jonah Benjamin Gadasi, Barbara Gwynne Simpson, Pedro Lomónaco, and Jesús María Blanco Ilzarbe. Impact of Limited Degree of Freedom Drag Coefficients on a Floating Offshore Wind Turbine Simulation. *Journal of Marine Science and Engineering*, 11, 2023.
- [231] Arne Alping and Gunnar Jacobsen. *Effects of Clipping and Limiting on the Performance of Optical Fiber Communication Systems*, volume 32. 2014.
- [232] Udo Zölzer. *AD/DA Conversion*. John Wiley & Sons, Ltd, 2011.
- [233] IEC. Wind energy generation systems – part 3-1: Design requirements for fixed offshore wind turbines, 2019.
- [234] Matthew J. Churchfield, Sang Lee, Patrick J. Moriarty, Luis A. Martinez, Stefano Leonardi, Ganesh Vijayakumar, and James G. Brasseur. A large-eddy simulation of wind-plant aerodynamics. Technical Report NREL/CP-500-53554, National Renewable Energy Laboratory, 2012. 50th AIAA Aerospace Sciences Meeting.

Bibliography

- [235] J. M. Jonkman, E. S. P. Branlard, and J. P. Jasa. Influence of wind turbine design parameters on linearized physics-based models in openfast. *Wind Energy Science*, 7(2):559–571, March 2022. NREL/JA-5000-81481.
- [236] Torgeir Moan, Zhen Gao, Erin E. Bachynski, and Amir R. Nejad. Recent Advances in Integrated Response Analysis of Floating Wind Turbines in a Reliability Perspective. *Journal of Offshore Mechanics and Arctic Engineering*, 142(5):052002, 03 2020.
- [237] Katarzyna Patryniak, Maurizio Collu, and Andrea Coraddu. Multidisciplinary design analysis and optimisation frameworks for floating offshore wind turbines: State of the art. *Ocean Engineering*, 251:111002, 2022.
- [238] Erin Elizabeth Bachynski. *Design and Dynamic Analysis of Tension Leg Platform Wind Turbines*. PhD thesis, Norwegian University of Science and Technology, Trondheim, 2014.
- [239] Matthew Hall, Stein Housner, Daniel Zalkind, Pietro Bortolotti, David Ogden, and Garrett E. Barter. An Open-Source Frequency-Domain Model for Floating Wind Turbine Design Optimization. *Journal of Physics Conference Series*, 2265(4), 2022.
- [240] Jose Azcona, Xabier Munduate, Leo Gonzalez, and Tor Anders Nygaard. Experimental validation of a dynamic mooring lines code with tension and motion measurements of a submerged chain. *Ocean Engineering*, 129:415–427, 2017.
- [241] Giovanni A. Amaral, Jordi Mas-Soler, and Alexandre N. Simos. Assessing the influence of nonlinear mooring restoring forces in the optimization and design of fowt. *Journal of Physics: Conference Series*, 2875(1):012037, 2024.
- [242] Wenjie Zhong, Weiwen Zhao, Decheng Wan, and Yan Zhao. Comparison study on mooring line models for hydrodynamic performances of floating offshore wind turbines. *Ocean Engineering*, 296:117083, 2024.

Bibliography

- [243] Johannes Palm, Claes Eskilsson, and Lars Bergdahl. An hp-adaptive discontinuous Galerkin method for modelling snap loads in mooring cables. *Ocean Engineering*, 144:266–276, 2017.
- [244] Lin Chen, Biswajit Basu, and Søren R.K. Nielsen. A coupled finite difference mooring dynamics model for floating offshore wind turbine analysis. *Ocean Engineering*, 162:304–315, 2018.
- [245] Pau Trubat, Climent Molins, and Xavier Gironella. Quasi-dynamic mooring line model. *Ocean Engineering*, 238:109739, 2021.
- [246] Imanol Touzon, Vincenzo Nava, Beatriz de Miguel, and Victor Petuya. Frequency domain modelling of a coupled system of floating structure and mooring lines: An application to a wave energy converter. *Ocean Engineering*, 220:108498, 2021.
- [247] Serag-Eldin Abdelmoteleb and Erin E. Bachynski-Polić. A frequency-domain optimization procedure for catenary and semi-taut mooring systems of floating wind turbines. *Marine Structures*, 101:103768, 2025.
- [248] J. A. P. Aranha and M. O. Pinto. Dynamic tension in risers and mooring lines: an algebraic approximation for harmonic excitation. *Applied Ocean Research*, 23(2):63–81, 2001.
- [249] Shi He and Aijun Wang. Time and frequency domain dynamic analysis of offshore mooring. *Journal of Marine Science and Engineering*, 9(7):781, 2021.
- [250] Eguzkiñe Martinez-Puente, Ander Zarketa-Astigarraga, Manex Martinez-Agirre, Alaitz Zabala, Jon Ander Esnaola, Miguel Muñoz-Calvente, Iñigo Llavori, and Markel Penalba. Benchmarking of spectral methods for fatigue assessment of mooring systems and dynamic cables in offshore renewable energy technologies. *Ocean Engineering*, 308:118311, 2024.
- [251] Wengang Mao, Björn Nembach, Wuliu Tian, and Da Wu. A spectral fatigue assessment framework for floating offshore wind turbines. *Ships and Offshore Structures*, 20(3):306–315, 2025.

Bibliography

- [252] John D. Sørensen and Jens N. Sørensen. *Wind Energy Systems: Optimising Design and Construction for Safe and Reliable Operation*. Woodhead Publishing, 2010.
- [253] Yupeng Song, Biswajit Basu, Zili Zhang, John Dalsgaard Sørensen, Jie Li, and Jianbing Chen. Dynamic reliability analysis of a floating offshore wind turbine under wind-wave joint excitations via probability density evolution method. *Renewable Energy*, 168:991–1014, 2021.
- [254] Kun Xu, Min Zhang, Yanlin Shao, Zhen Gao, and Torgeir Moan. Effect of wave nonlinearity on fatigue damage and extreme responses of a semi-submersible floating wind turbine. *Applied Ocean Research*, 91:101879, 2019.
- [255] Shuaishuai Wang, Torgeir Moan, and Zhen Gao. Methodology for global structural load effect analysis of the semi-submersible hull of floating wind turbines under still water, wind, and wave loads. *Marine Structures*, 91, 2023.
- [256] Fengshen Zhu, Baran Yeter, Feargal Brennan, and Maurizio Collu. Time-domain fatigue reliability analysis for floating offshore wind turbine substructures using coupled nonlinear aero-hydro-servo-elastic simulations. *Engineering Structures*, 318:118759, 2024.
- [257] Turan Dirlik. Application of computer in fatigue analysis, 1985.
- [258] R. Tovo and D. Benasciutti. A simplified approach to evaluate fatigue damage under wide-band stationary random processes. *International Journal of Fatigue*, 26(9):831–839, 2004.
- [259] Turan Dirlik and Denis Benasciutti. Dirlik and tovo-benasciutti spectral methods in vibration fatigue: A review with a historical perspective. *Metals*, 11(9), 2021.
- [260] Nestor V. Queipo, Raphael T. Haftka, Wei Shyy, Tushar Goel, Rajkumar Vaidyanathan, and P. Kevin Tucker. Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41(1):1–28, 2005.

Bibliography

- [261] Christopher M. Bishop and Hugh Bishop. *Deep Learning: Foundations and Concepts*. Springer, Cham, 2023.
- [262] Luca Oneto. *Model Selection and Error Estimation in a Nutshell*, volume 15 of *Modeling and Optimization in Science and Technologies*. Springer, Cham, 2020.
- [263] A. Pegalajar-Jurado, M. Borg, and H. Bredmose. An efficient frequency-domain model for quick load analysis of floating offshore wind turbines. *Wind Energy Science*, 3(2):693–712, 2018.
- [264] J. Jonkman, S. Butterfield, W. Musial, and G. Scott. Definition of a 5-MW Reference Wind Turbine for Offshore System Development, 2009.
- [265] S A Ning. CCBlade Documentation. Release 0.1.0, 2013.
- [266] Openfast v2.6.
- [267] DNV. Offshore standard dnv-os-e301. position mooring, 2013. DNV.
- [268] DNVGL-RP-C203: Fatigue Design of Offshore Steel Structures. Technical report, DNV GL, 2016. Recommended Practice.
- [269] Turan Dirlik and Denis Benasciutti. Dirlik and tovo-benasciutti spectral methods in vibration fatigue: A review with a historical perspective. *Metals*, 11(9), 2021.
- [270] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [271] S. P. Adam, S. A. N. Alexandropoulos, P. M. Pardalos, and M. N. Vrahatis. No free lunch theorem: A review. *Approximation and optimization*, pages 57–82, 2019.
- [272] J. Shawe-Taylor and N. Cristianini. *Kernel methods for pattern analysis*. Cambridge university press, 2004.
- [273] Z. H. Zhou. *Ensemble methods: foundations and algorithms*. CRC press, 2025.

Bibliography

- [274] C. Silvano, D. Ielmini, F. Ferrandi, L. Fiorin, and Others. A survey on deep learning hardware accelerators for heterogeneous hpc platforms. *ACM Computing Surveys*, 57(11):1–39, 2025.
- [275] M. Z. Naser and A. H. Alavi. Error metrics and performance fitness indicators for artificial intelligence and machine learning in engineering and sciences. *Architecture, Structures and Construction*, 3(4):499–517, 2023.
- [276] K. L. Sainani. The value of scatter plots. *PM&R*, 8(12):1213–1217, 2016.
- [277] NREL. Aerodyn users guide and theory manual. <https://openfast.readthedocs.io/en/main/source/user/aerodyn/index.html>, 2026.
- [278] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [279] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [280] S. Bai, J. Z. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- [281] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, and Others. Attention is all you need. *Neural information processing systems*, 2017.
- [282] A. Das, W. Kong, R. Sen, and Y. Zhou. A decoder-only foundation model for time-series forecasting. In *International Conference on Machine Learning*, 2024.
- [283] MDO Lab. Ipopt — pyoptsparse documentation. Online; accessed 7 Nov 2025, 2025. <https://mdolab-pyoptsparse.readthedocs-hosted.com/en/latest/optimizers/IPOPT.html>.
- [284] Ye An, Zhisheng Xia, Min Luo, Xin Jin, and Ronghua Zhu. Surrogate model-assisted multi-objective optimization of the mooring system of a floating wind

Bibliography

- turbine considering short- and long-term performances. *Energy*, page 139610, 2025.
- [285] Serag Abdulmotaleb and Matthew Hall. Add: Line dynamics + new compositeline class. <https://github.com/NREL/MoorPy/commit/45d2827863b4067c2f0195c9dcf992fcb44a78d0>, June 2023. GitHub, commit 45d2827.
- [286] Kapil Khanal, Carlos A. Michelén Ströfer, Matthieu Ancellin, and Maha N. Haji. Fully differentiable boundary element solver for hydrodynamic sensitivity analysis of wave-structure interactions. *Applied Ocean Research*, 163:104707, October 2025.
- [287] N. J. Abbas, D. S. Zalkind, L. Pao, and A. Wright. A reference open-source controller for fixed and floating offshore wind turbines. *Wind Energy Science*, 7(1):53–73, 2022.
- [288] National Renewable Energy Laboratory (NREL). User’s guide to bmodes (software for computing rotating beam-coupled modes). Technical Report NREL/TP-500-39133, National Renewable Energy Laboratory, December 2005. Technical Report.
- [289] G. K. V. Ramachandran, A. Robertson, J. M. Jonkman, and M. D. Masciola. Investigation of response amplitude operators for floating offshore wind turbines. In *Proceedings of the 23rd International Ocean, Offshore and Polar Engineering Conference (ISOPE 2013)*, Anchorage, Alaska, USA, June 30–July 5 2013. ISOPE. Preprint, NREL/CP-5000-58098.
- [290] *Eurocode 3: Design of steel structures – Part 1-1: General rules and rules for buildings*. European Committee for Standardization (CEN), Brussels, 2005. EN 1993-1-1.

Appendix A

Aerodynamic damping

The motions of the floating platform affect the velocity of the inflow seen by the rotor, hence influencing the aerodynamic loads. The aerodynamic damping, defined as the change of aerodynamic load due to the change of inflow velocity, can be obtained numerically by varying the inflow velocity for a stationary rotor and observing the change of the resulting aerodynamic loads:

$$B_{aero} = \frac{dF_{aero}}{dv} \quad (\text{A.1})$$

In general, there are three aerodynamic forces, three moments, and three inflow directions. However, the vertical and lateral motions and loads are significantly lower than the horizontal motions, roughly aligned with the rotor axis. Therefore, the majority of the terms in the aerodynamic damping matrix (or aerodynamic load Jacobian) are negligible.

Aerodynamic load depends not only on the inflow velocity, but also on the inflow direction and control action. For this reason, the damping is computed for different values of the mean platform inclination angles (which are equivalent to the rotor plane inclination angles assuming a rigid structure). For each simulation, the control variables (rotor speed and blade pitch) are kept fixed at the values from the control table.

The baseline model described in Chapter 3 uses aerodynamic damping precomputed for a range of inflow speeds and inclination angles using a time domain Blade-Element

Appendix A. Aerodynamic damping

Momentum (BEM) code AeroDyn [188] within OpenFAST. The more involved model of Chapter 6 uses direct computation of aerodynamic damping using a quasi-static BEM code CCBlade.

A.1 Precomputed look-up tables

The precomputation was done for a grid of wind speeds sampled every 1 ms^{-1} and inclination angles sampled at 0.5° . The central difference numerical derivative at a given grid point was obtained based on a small perturbation of the wind speed δv :

$$B_{aero}(v_i, \theta_i) = \frac{F_{aero}(v_i + \delta v, \theta_i) - F_{aero}(v_i - \delta v, \theta_i)}{2\delta v} \quad (\text{A.2})$$

The simulations were run in steady wind. Only the surge-surge component of the aerodynamic damping matrix was computed, and the surge-pitch and pitch-pitch terms were estimated from it following the approach in [195] based on the moment the aerodynamic force creates at the origin z_h below the hub height:

$$B_{aero,(1,5)} = B_{aero,(5,1)} = B_{aero,(1,1)}z_h \quad (\text{A.3})$$

$$B_{aero,(5,5)} = B_{aero,(1,1)}z_h^2 \quad (\text{A.4})$$

The coefficients obtained this way were stored in look-up tables and interpolated using a linear interpolation when required by the overall framework.

A.2 Direct aerodynamic damping

The precomputation method relied on the steady-state results of aerodynamic load resulting from a steady wind inflow. Therefore, it was possible to avoid the uncertainties due to interpolation and directly obtain the aerodynamic damping using a quasi-static BEM model CCBlade [265].

Again, subject to the mean inclination and wind speed input, and fixed control variables, CCBlade calculates the derivatives of the components of aerodynamic load

Appendix A. Aerodynamic damping

with respect to the inflow in the x -direction:

$$B_{aero,(1,1)} = \frac{\delta F_{aero,x}}{\delta U} \quad (\text{A.5})$$

$$B_{aero,(2,1)} = \frac{\delta F_{aero,y}}{\delta U} \quad (\text{A.6})$$

$$B_{aero,(3,1)} = \frac{\delta F_{aero,z}}{\delta U} \quad (\text{A.7})$$

These are then used to compute the remaining items of the aerodynamic damping matrix like before in Equations A.3–A.4.

The aerodynamic force and damping curves obtained this way are presented in Figure A.1 together with the control parameters (blade pitch and rotor speed).

Appendix A. Aerodynamic damping




Figure A.1: Aerodynamic force (thrust) and damping in surge direction curves for three values of mean pitch angle $\bar{\eta}_\theta$. The blade pitch (θ) and rotor speed (Ω) curves are shown for reference.

Appendix A

Nonlinear solvers algorithms and diagrams

This appendix details the iterative solves, complementing Chapter 3 with algorithms and diagrams. Please refer to the chapter for nomenclature.

A.1 Nonlinear Block Jacobi solver

Algorithm 1 Nonlinear Block Jacobi Solver (NLBJ)

Inputs:

$$\hat{u}^{(0)} = [\hat{u}_1^{(0)}, \dots, \hat{u}_m^{(0)}]$$

Outputs:

$$\hat{u} = [\hat{u}_1, \dots, \hat{u}_m]$$

$$k = 0$$

While $\|\hat{u}^{(k)} - \hat{u}^{(k-1)}\|_2 > \varepsilon$ or $k = 0$:

For all each $i \in \{1, \dots, m\}$:

$$\hat{u}_i^{(k+1)} \leftarrow \text{solve } r_i \left(\hat{u}_i^{(k+1)}; \hat{u}_j^{(k)} \right) = 0, \quad j \neq i$$

$$k = k + 1$$

Algorithm 2 Nonlinear Block Gauss-Seidel Algorithm with Aitken Acceleration

Inputs:

$$\hat{u}^{(0)} = [\hat{u}_1^{(0)}, \dots, \hat{u}_m^{(0)}]$$

$$\theta^{(0)}$$

Outputs:

$$\hat{u} = [\hat{u}_1, \dots, \hat{u}_m]$$

$$k = 0$$

While $\|\hat{u}^{(k)} - \hat{u}^{(k-1)}\|_2 > \varepsilon$ **or** $k = 0$ **do**:

For each $i \in \{1, \dots, m\}$ **do**:

$$\hat{u}_{\text{temp}} \leftarrow \text{solve } r_i \left(\hat{u}_i^{(k+1)}; \hat{u}_1^{(k+1)}, \dots, \hat{u}_{i-1}^{(k+1)}, \hat{u}_{i+1}^{(k)}, \dots, \hat{u}_m^{(k)} \right) = 0$$

$$\Delta \hat{u}_i^{(k)} = \hat{u}_{\text{temp}} - \hat{u}_i^{(k)}$$

If $k > 0$ **then**:

$$\theta^{(k)} = \theta^{(k-1)} \left(1 - \frac{(\Delta \hat{u}^{(k)} - \Delta \hat{u}^{(k-1)})^\top \Delta \hat{u}^{(k)}}{\|\Delta \hat{u}^{(k)} - \Delta \hat{u}^{(k-1)}\|^2} \right)$$

end if

$$\hat{u}_i^{(k+1)} = \hat{u}_i^{(k)} + \theta^{(k)} \Delta \hat{u}_i^{(k)}$$

end for

$$k = k + 1$$

end while

Appendix A. Nonlinear solvers algorithms and diagrams




Figure A.1: Nonlinear Block Jacobi solver. Reproduced from [11].

A.2 Nonlinear Block Gauss-Seidel solver




Figure A.2: Nonlinear Block Gauss-Seidel solver. Reproduced from [11].

A.3 Monolithic Newton's solver




Figure A.3: Full-space hierarchical Newton solver. Reproduced from [11].

Algorithm 3 Full-space hierarchical Newton Algorithm

Inputs:

$$u^{(0)} = [u_1^{(0)}, \dots, u_n^{(0)}]$$

Outputs:

$$u = [u_1, \dots, u_n]$$

$$k = 1$$

while $\|r\|_2 > \varepsilon$ **do:**

for all $i \in \{1, \dots, n\}$ **do:**

while $\|r_i\|_2 > \varepsilon$ **do:**

 compute $r_i(u_i^{(k)}; u_{j \neq i}^{(k-1)})$

 compute $\frac{\partial r_i}{\partial u_i}$

 solve $\frac{\partial r_i}{\partial u_i} \Delta u_i = -r_i$

$$u_i^{(k)} = u_i^{(k-1)} + \Delta u_i$$

end while

end for

 compute $r(u^{(k)})$

 compute $\frac{\partial r}{\partial u}$

 solve $\frac{\partial r}{\partial u} \Delta u = -r$

$$u^{(k+1)} = u^{(k)} + \Delta u$$

$$k = k + 1$$

end while

Appendix B

Nonlinear solvers examples

This appendix presents an example calculation demonstrating the different nonlinear solvers trialled in Chapter 3.

Starting with a 1D problem, we want to solve the implicit nonlinear equation

$$x = \cos x \quad \Longleftrightarrow \quad F(x) = \cos x - x = 0,$$

starting from $x_0 = 0.5$. The true solution is $x^* \approx 0.739085133$.

1. Newton's method

Derivation

Given $F : \mathbb{R} \rightarrow \mathbb{R}$ and a current iterate x_k , form the first-order Taylor expansion

$$F(x_k + s) \approx F(x_k) + F'(x_k) s.$$

Newton chooses the step s_k so that the *linearised* equation is solved exactly:

$$F(x_k) + F'(x_k) s_k = 0 \quad \Longrightarrow \quad s_k = -\frac{F(x_k)}{F'(x_k)}.$$

Appendix B. Nonlinear solvers examples

Hence the iteration

$$x_{k+1} = x_k + s_k = x_k - \frac{F(x_k)}{F'(x_k)}.$$

Applied to $F(x) = \cos x - x$

Here $F'(x) = -\sin x - 1$, so

$$x_{k+1} = x_k - \frac{\cos x_k - x_k}{-\sin x_k - 1}.$$

Starting $x_0 = 0.5$:

$$x_1 = 0.755222417, \quad x_2 = 0.739141666, \quad x_3 = 0.7390851339, \quad x_4 = 0.7390851332.$$

(Quadratic convergence near the root.)

2. Broyden / Secant

Derivation (Secant in 1D)

Newton uses the true derivative $F'(x_k)$. The *secant* method replaces it by the finite-difference slope from the two latest points:

$$F'(x_k) \approx \frac{F(x_k) - F(x_{k-1})}{x_k - x_{k-1}}.$$

Plugging this into the Newton step formula gives

$$s_k = -\frac{F(x_k)}{F'(x_k)} \approx -F(x_k) \frac{x_k - x_{k-1}}{F(x_k) - F(x_{k-1})},$$

and thus the iteration

$$x_{k+1} = x_k - F(x_k) \frac{x_k - x_{k-1}}{F(x_k) - F(x_{k-1})}.$$

In multiple dimensions, Broyden generalises this idea by maintaining and low-rank updating a Jacobian approximation B_k so that $B_{k+1}s_k \approx y_k$ where $s_k = x_{k+1} - x_k$ and

Appendix B. Nonlinear solvers examples

$y_k = F(x_{k+1}) - F(x_k)$; in 1D this collapses exactly to the secant update above.

Applied to $F(x) = \cos x - x$

With $x_0 = 0.5$, $x_1 = 1.0$:

$$x_2 = 0.725481587, \ x_3 = 0.738398620, \ x_4 = 0.739087210, \ x_5 = 0.7390851329.$$

(Superlinear convergence; no derivatives needed.)

3. Fixed-point iteration: deriving the step

Derivation (from a rearrangement)

Rewrite $F(x) = 0$ as a fixed-point equation $x = G(x)$. A generic fixed-point method sets

$$x_{k+1} = G(x_k)$$

which can be seen as solving $x - G(x) = 0$ via the trivial linearization $x_{k+1} - G(x_k) = 0$. Convergence is guaranteed, e.g. by the Banach fixed-point theorem, if G is a contraction near the solution, i.e. $|G'(x^*)| < 1$.

Applied as Jacobi/Gauss–Seidel in 1D

Take $G(x) = \cos x$:

$$x_{k+1} = \cos x_k.$$

For a *scalar* problem, “Jacobi” and “Gauss–Seidel” are identical because there is only one component to update.

Starting $x_0 = 0.5$:

$$x_1 = 0.877582562, \ x_2 = 0.639012494, \ x_3 = 0.802685101, \ x_4 = 0.694778027, \ x_5 = 0.768195831, \\ x_{10} = 0.739006780, \ x_{15} = 0.739085131, \ x_{20} = 0.739085133.$$

Here $|G'(x^*)| = |-\sin x^*| \approx 0.673 < 1$, so convergence is linear.

4. How Jacobi vs. Gauss–Seidel differ in systems

For a system $F(x) = 0$ with $x = (x_1, \dots, x_n)$, a common route is to write a component-wise fixed-point form

$$x_i = G_i(x_1, \dots, x_n), \quad i = 1, \dots, n.$$

Jacobi forms $x^{(k+1)}$ using only the previous iterate on the right-hand side:

$$x_i^{(k+1)} = G_i(x_1^{(k)}, \dots, x_n^{(k)}).$$

Gauss–Seidel immediately reuses the latest components as they are computed (lexicographic update):

$$x_i^{(k+1)} = G_i(x_1^{(k+1)}, \dots, x_{i-1}^{(k+1)}, x_i^{(k)}, \dots, x_n^{(k)}).$$

Thus, the *derived step* for GS differs by plugging in the freshest available values on the right-hand side, which often accelerates convergence but sacrifices parallelism.

B.1 2D example

Consider the nonlinear system:

$$\begin{cases} y_1 = \cos(y_2), \\ y_2 = \sin(y_1), \end{cases} \iff \begin{cases} F_1(y_1, y_2) = y_1 - \cos(y_2) = 0, \\ F_2(y_1, y_2) = y_2 - \sin(y_1) = 0. \end{cases}$$

We start from the initial guess $y_1^{(0)} = 0.5$, $y_2^{(0)} = 0.5$.

B.1.1 Jacobi iteration

Update both components using the previous iteration values:

$$y_1^{(k+1)} = \cos(y_2^{(k)}), \quad y_2^{(k+1)} = \sin(y_1^{(k)}).$$

Appendix B. Nonlinear solvers examples

Iterations:

$$\begin{aligned}y_1^{(1)} &= \cos(0.5) \approx 0.8776, & y_2^{(1)} &= \sin(0.5) \approx 0.4794, \\y_1^{(2)} &= \cos(0.4794) \approx 0.8872, & y_2^{(2)} &= \sin(0.8776) \approx 0.7695.\end{aligned}$$

Convergence is linear, as only previous values are used.

B.1.2 Gauss–Seidel iteration

Update sequentially using the freshest available values:

$$y_1^{(k+1)} = \cos(y_2^{(k)}), \quad y_2^{(k+1)} = \sin(y_1^{(k+1)}).$$

Iterations:

$$\begin{aligned}y_1^{(1)} &= \cos(0.5) \approx 0.8776, & y_2^{(1)} &= \sin(0.8776) \approx 0.7695, \\y_1^{(2)} &= \cos(0.7695) \approx 0.7180, & y_2^{(2)} &= \sin(0.7180) \approx 0.6571.\end{aligned}$$

Convergence is faster than Jacobi because updated values propagate immediately.

B.1.3 Hierarchical full-space Newton

Define the residual vector and Jacobian:

$$R(y_1, y_2) = \begin{bmatrix} y_1 - \cos(y_2) \\ y_2 - \sin(y_1) \end{bmatrix}, \quad J = \begin{bmatrix} 1 & \sin(y_2) \\ -\cos(y_1) & 1 \end{bmatrix}.$$

Iteration procedure:

1. Perform a few inner sweeps to reduce local residuals (hierarchical step).
2. Assemble full residual $R(y^{(k)})$ and Jacobian $J(y^{(k)})$.
3. Solve the linear system $J\Delta y = -R$ for the Newton step Δy .
4. Update $y^{(k+1)} = y^{(k)} + \Delta y$.

Appendix B. Nonlinear solvers examples

Step 0: $y^{(0)} = (0.5, 0.5)$ Inner sweeps (hierarchical preconditioning): perform one Gauss–Seidel style sweep on the components to reduce local residuals:

$$\tilde{y}_1 = \cos(y_2^{(0)}) \approx 0.8776, \quad \tilde{y}_2 = \sin(\tilde{y}_1) \approx 0.7695$$

These values $\tilde{y}$ provide a better starting point for the global Newton step. Assemble full residual and Jacobian at $\tilde{y}$:

$$R(\tilde{y}) \approx \begin{bmatrix} 0.8776 - \cos(0.7695) \\ 0.7695 - \sin(0.8776) \end{bmatrix} \approx \begin{bmatrix} 0.1568 \\ 0.0000 \end{bmatrix}, \quad J(\tilde{y}) \approx \begin{bmatrix} 1 & 0.6951 \\ -0.6361 & 1 \end{bmatrix}.$$

Solve linear system: $J\Delta y = -R$ to obtain Newton step:

$$\Delta y \approx \begin{bmatrix} -0.1473 \\ -0.0890 \end{bmatrix}.$$

Update state:

$$y^{(1)} = \tilde{y} + \Delta y \approx \begin{bmatrix} 0.7303 \\ 0.6805 \end{bmatrix}.$$

Next iteration: repeat inner sweeps and full-space Newton until $\|R(y^{(k)})\|$ is below tolerance. Quadratic convergence is observed after a few iterations.

Appendix C

Flexible tower and modal analysis

This appendix details the flexible tower properties and modal analysis as required to evaluate the natural frequencies of the FOWT, supplementing the resonance avoidance constraints defined in Chapter 2.2.1.

The tower is modelled as a vertical cantilever beam, fixed at its base and free at its top. The structure is assumed to be isotropic and prismatic along its height, with geometry and material properties defined at discrete stations. The outer radius and wall thickness (*effective thickness*, accounting for the flanges, welds, flanges, and paint not accounted for otherwise) are prescribed at the tower base and top, and interpolated linearly between these points. For a tower section with outer radius r_o , wall thickness

Appendix C. Flexible tower and modal analysis

t , and inner radius $r_i = r_o - t$, the following sectional properties are defined:

$$A = \pi (r_o^2 - r_i^2) \quad \text{Cross-sectional area} \quad (\text{C.1})$$

$$m = \rho A \quad \text{Mass per unit length} \quad (\text{C.2})$$

$$I_{xx} = I_{yy} = \frac{\pi}{4} (r_o^4 - r_i^4) \quad \text{Area moment of inertia} \quad (\text{C.3})$$

$$I_{xx,\text{mass}} = m \frac{I_{xx}}{A} \quad \text{Mass moment of inertia per unit length} \quad (\text{C.4})$$

$$S_{\text{bend}} = E I_{xx} \quad \text{Bending stiffness} \quad (\text{C.5})$$

$$J = \frac{\pi}{2} (r_o^4 - r_i^4) \quad \text{Polar moment of inertia} \quad (\text{C.6})$$

$$S_{\text{torsion}} = G J \quad \text{Torsional stiffness} \quad (\text{C.7})$$

$$S_{\text{axial}} = E A \quad \text{Axial stiffness} \quad (\text{C.8})$$

Here ρ is the material density, E is Young's modulus, and G is the shear modulus.

These properties are evaluated at a set of equally spaced tower stations, providing a discretised representation of the tower. Results for the reference tower are listed in Table C.1. The rotor–nacelle assembly is treated as a lumped mass at the tower top, with its equivalent mass, centre of gravity, and inertia tensor derived from the contributions of the nacelle and rotor. The floating platform is represented as a rigid body with specified mass, centre of gravity, and mass moments of inertia. The hydrodynamic added-mass, hydrostatic restoring stiffness, and mooring stiffness matrices are incorporated into the system to account for the floating foundation.

The natural frequencies of the system are obtained using the finite element code BModes developed at NREL [288]. The tower is represented as a Euler–Bernoulli beam with distributed stiffness and mass properties along its span, and the coupled non-linear governing equations are derived from Hamilton's principle. These equations are analytically linearised and discretised using a special-purpose finite element with 15 degrees of freedom per element, enabling the representation of coupled bending, torsional, and axial deformations. Shear deformation and cross-sectional warping are neglected, consistent with the Euler–Bernoulli beam theory.

From the element formulations and boundary conditions, BModes assembles the

Appendix C. Flexible tower and modal analysis

Table C.1: Tower sectional properties of the reference system evaluated at 11 stations.

z [m]	r [m]	t [m]	A [m ²]	m [kg m ⁻¹]	$I_{xx,\text{mass}}$ [kg m]	EI [GN m ²]	GJ [GN m ²]	EA [GN]
10.00	3.25	0.03	0.55	4667.0	24443.7	604	465	1150
17.76	3.12	0.03	0.51	4345.3	20952.2	518	398	1070
25.52	2.99	0.03	0.47	4034.8	17847.0	441	339	1000
33.28	2.86	0.02	0.44	3735.4	15098.5	373	287	923
41.04	2.72	0.02	0.41	3447.3	12678.6	313	241	852
48.80	2.59	0.02	0.37	3170.4	10560.1	261	201	783
56.56	2.46	0.02	0.34	2904.7	8717.2	215	166	718
64.32	2.33	0.02	0.31	2650.2	7125.0	176	135	655
72.08	2.20	0.02	0.28	2406.9	5759.8	142	110	595
79.84	2.07	0.02	0.26	2174.8	4599.3	114	87	537
87.60	1.94	0.02	0.23	1953.9	3622.1	89.5	68.9	483

global mass and stiffness matrices of the coupled tower–platform system (ignoring gyroscopic effect). The free vibration of the coupled tower–platform system is governed by the second-order equations of motion

$$\mathbf{M}\ddot{q}(t) + \mathbf{K}q(t) = \mathbf{0}, \quad (\text{C.9})$$

where $q(t)$ is the vector of generalised displacements, $\dot{q}(t)$ and $\ddot{q}(t)$ are the corresponding velocities and accelerations, and $\mathbf{M}$ and $\mathbf{K}$ are the global mass and stiffness matrices. Assuming harmonic motion of the form

$$q(t) = \phi e^{i\omega t}, \quad (\text{C.10})$$

where ϕ is the mode shape and ω the natural circular frequency, substitution gives the quadratic eigenvalue problem

$$(-\omega^2\mathbf{M} + \mathbf{K})\phi = \mathbf{0}. \quad (\text{C.11})$$

Rather than solving this quadratic problem directly, BModes introduces a state vector that stacks displacements and velocities,

$$x(t) = \begin{bmatrix} q(t) \\ \dot{q}(t) \end{bmatrix}, \quad \dot{x}(t) = \begin{bmatrix} \dot{q}(t) \\ \ddot{q}(t) \end{bmatrix}. \quad (\text{C.12})$$

Appendix C. Flexible tower and modal analysis

Using the equations of motion $\ddot{q}(t) = -\mathbf{M}^{-1}\mathbf{K}q(t)$, the system becomes

$$\dot{x}(t) = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & \mathbf{0} \end{bmatrix}}_{\mathbf{A}} x(t). \quad (\text{C.13})$$

The resulting first-order system can be written as the linear eigenvalue problem

$$\mathbf{A}x = \lambda x, \quad (\text{C.14})$$

with eigenvalues $\lambda = i\omega$. The natural frequencies are finally obtained as

$$f = \frac{\omega}{2\pi}. \quad (\text{C.15})$$

Solving this eigenvalue problem yields the natural frequencies and associated mode shapes. Note that the mode shapes and frequencies are not prescribed a priori but are obtained directly from the eigenvalue problem. Among the resulting modes, some correspond predominantly to rigid-body motions of the floating platform (surge, heave, pitch, which are coupled through hydrostatic and mooring restoring), while others represent flexible tower modes (fore–aft and side–side bending). The natural frequencies of the reference floating wind turbine system corresponding to the surge, heave, pitch, and tower 1st fore–aft modes are summarised in Table C.2. These were verified against the values reported in [289].

Table C.2: Natural frequencies of the reference floating wind turbine system.

Mode No.	Frequency [Hz]	Dominant motion
1	0.008	Surge (rigid-body, coupled)
2	0.032	Heave (rigid-body, coupled)
3	0.034	Pitch (rigid-body, coupled)
4	0.4778	Tower fore–aft, 1st bending

Appendix D

Tower structural checks

This Appendix provides further details of the buckling and yielding checks (constraints) introduced in Chapter 4.

D.1 Shell (local) buckling

The tower is modelled as a circular, linearly tapered in radius and thickness, conical shell without longitudinal stiffeners.¹ The taper angle is assumed to be small in the order of 1° (for perspective, the original tower's slope is about 0.8°). It is subjected to a design axial force N_{Sd} and a global bending moment M_{Sd} arising from aerodynamic load on the rotor and tower, self-weight and inertial effects. In the typical ultimate limit state conditions, hoop stress from external pressure and in-plane shear are small compared to longitudinal compression and are neglected in the shell-buckling check.

The conical shell of length L is mapped to an equivalent cylinder, following the recommendations of DNV-RP-C202 [212], Sections 4.3–4.4, by defining the equivalent radius, thickness, and length:

$$\begin{aligned}\alpha &= \arctan\left(\frac{R_{\text{base}} - R_{\text{top}}}{L}\right), & \bar{R} &= \frac{1}{2}(R_{\text{top}} + R_{\text{base}}), & \bar{t} &= \frac{1}{2}(t_{\text{top}} + t_{\text{base}}), \\ R_e &= \frac{\bar{R}}{\cos \alpha}, & t_e &= \bar{t} \cos \alpha, & l_e &= \frac{L}{\cos \alpha}.\end{aligned}$$

¹The transverse stiffeners/flanges are accounted for by using the *effective* material properties and wall thickness.

Appendix D. Tower structural checks

Given the equivalent particulars of the tower and the design resultant loads (N_{Sd} , M_{Sd}), the resultant longitudinal stresses at the extreme location are given by:

$$\begin{aligned}\sigma_{a,Sd} &= \frac{N_{Sd}}{2\pi R_e t_e} \quad (\text{uniform axial membrane stress}), \\ \sigma_{m,Sd} &= \frac{|M_{Sd}|}{\pi R_e^2 t_e} \quad (\text{bending stress at extreme location}),\end{aligned}$$

and the total compressive stress is obtained as:

$$\sigma_{x,Sd}^{\max} = \sigma_{a,Sd} + \sigma_{m,Sd}.$$

Following Table 3-2 and Sections 3.4 of DNV-RP-C202, buckling coefficients are computed to account for the buckling susceptibility and real design conditions, to effectively reduce the “perfect cylinder” buckling strength to a realistic one:

$$\begin{aligned}Z_\ell &= \frac{l_e^2}{R_e t_e} \sqrt{1 - \nu^2} \quad (\text{Curvature parameter}), \\ \xi &= 0.702 Z_\ell \quad (\text{Buckling coefficient}), \\ \rho_{\text{axial}} &= \frac{1}{2} \left(1 + \frac{R_e}{150 t_e} \right)^{-1/2} \quad (\text{Imperfection sensitivity reduction factor for axial load}), \\ \rho_{\text{bending}} &= \frac{1}{2} \left(1 + \frac{R_e}{300 t_e} \right)^{-1/2} \quad (\text{Imperfection sensitivity reduction factor for bending}), \\ C_{\text{Axial}} &= \Xi \sqrt{1 + \left(\frac{\rho_{\text{axial}} \xi}{\Xi} \right)^2} \quad (\text{Reduced buckling coefficient for axial load}), \\ C_{\text{Bending}} &= \Xi \sqrt{1 + \left(\frac{\rho_{\text{bending}} \xi}{\Xi} \right)^2} \quad (\text{Reduced buckling coefficient for bending}).\end{aligned}$$

The Ξ coefficient accounts for fabrication quality/imperfection class. For unstiffened cylindrical shells, it can be taken as 1.0 (default imperfection class). The corresponding elastic buckling strengths are then obtained by reducing the ideal elastic buckling stress by the factors above:

$$f_{Ea} = C_{\text{Axial}} \frac{\pi^2 E}{12(1 - \nu^2)} \left(\frac{t_e}{l_e} \right)^2, \quad f_{Em} = C_{\text{Bending}} \frac{\pi^2 E}{12(1 - \nu^2)} \left(\frac{t_e}{l_e} \right)^2.$$

Appendix D. Tower structural checks

The characteristic buckling strength f_{ks} is neither yield strength nor just elastic buckling strength, but a blend of the two depending on slenderness. The failure of the thick-walled shells will be primarily governed by material yielding, while the thin-walled structures will buckle before undergoing yielding. The relative importance of these two effects is modelled through the reduced shell slenderness factor:

$$\Lambda_S^2 = \frac{f_y}{\sigma_{x,Sd}^{\max}} \left(\frac{\sigma_{a,Sd}}{f_{Ea}} + \frac{\sigma_{m,Sd}}{f_{Em}} \right).$$

f_y is the characteristic yield strength of the tower material. For low slenderness ($\Lambda_S \ll 1$), $f_{ks} \approx f_y$ (yield-controlled), while for high slenderness, f_{ks} approaches the elastic buckling strength. This smooth transition, defined by DNV-RP-C202, is expressed as

$$f_{ks} = \frac{f_y}{\sqrt{1 + \Lambda_S^4}}.$$

The ratio f_{ks}/f_y is plotted against Λ_S in Figure D.1, illustrating how the characteristic buckling strength reduces from the yield strength at low slenderness to the buckling-controlled regime at higher slenderness.




Figure D.1: Transition from yield to buckling strength.

Reducing characteristic buckling strength by the material factor (here, taken equal to 1.1) gives the design buckling strength:

$$f_{kd} = \frac{f_{ks}}{\gamma_M}.$$

Appendix D. Tower structural checks

Finally, the utilisation for shell buckling is

$$U_{\text{shell}} = \frac{\sigma_{x,\text{Sd}}^{\max}}{f_{kd}} \leq 1.0.$$

as seen in Chapter 4.

D.2 Column (global) buckling

In addition to local shell buckling, the tower as a whole can fail by global column buckling, i.e. lateral instability of the entire tower treated as a cantilever beam. The tower of total height L is modelled as a cantilever column, i.e. fixed at the base and free at the top, with effective length factor $K = 2.0$ (standard for cantilever columns). For the column check, the section properties are evaluated at the tower base, where the axial force and bending moment are largest. The cross-sectional area is denoted by A , the second moment of area about the horizontal axis by I , and the section modulus by

$$W = \frac{I}{R_{\text{outer}}},$$

where R_{outer} is the outer radius of the base section. With these definitions, the Euler critical axial load is

$$N_{\text{cr}} = \frac{\pi^2 EI}{(KL)^2}.$$

A non-dimensional slenderness parameter is then defined by

$$\bar{\lambda} = \sqrt{\frac{Af_y}{N_{\text{cr}}}},$$

The quantity Af_y is the so-called *squash load*, i.e. the maximum axial force that the cross-section could sustain if the entire wall reached the yield strength f_y in uniform compression.

Similar to the shell buckling process, to account for imperfections and residual stresses, Eurocode 3 introduces a reduction factor χ , obtained via the Perry–Robertson

Appendix D. Tower structural checks

formula:

$$\phi = \frac{1}{2} \left[1 + \alpha(\bar{\lambda} - 0.2) + \bar{\lambda}^2 \right],$$

$$\chi = \frac{1}{\phi + \sqrt{\phi^2 - \bar{\lambda}^2}},$$

where α is the imperfection factor associated with the chosen buckling curve. Buckling curves are empirical reductions to the ideal Euler strength to account for initial out-of-straightness, residual stresses, and other imperfections that reduce real column strength. Different curves (presented in Figure D.2) correspond to different sensitivity levels. Typically $\alpha \in [0.21, 0.49]$, depending on fabrication quality. In this thesis, $\alpha = 0.34$ is adopted, corresponding to buckling curve b, which is typical for welded circular cylinders.




Figure D.2: EN 1993-1-1 Eurocode 3 buckling curves.

The design axial and bending resistances of the column are then given by:

$$N_{Rd} = \frac{\chi A f_y}{\gamma_M}, \quad M_{Rd} = \frac{W f_y}{\gamma_M}.$$

Finally, the column utilisation is checked using a linear interaction between axial load and bending moment (which is a conservative approach):

$$U_{\text{col}} = \frac{N_{\text{Sd}}}{N_{Rd}} + \frac{|M_{\text{Sd}}|}{M_{Rd}} \leq 1.0,$$

as formulated in Chapter 4.

D.3 Yielding check

In addition to stability checks against local shell buckling and global column buckling, the tower sections must also be verified against material yielding in accordance with Eurocode 3 (EN 1993-1-1) and DNV-RP-C202 [212, 290]. The design von Mises equivalent stress is evaluated as

$$\sigma_{\text{VM,Sd}} = \sqrt{(\sigma_{az} + \sigma_b)^2 + 3\tau^2},$$

where σ_{az} is the axial membrane stress due to the total axial force, σ_b is the bending stress at the extreme location, and τ is the in-plane shear stress from horizontal shear forces. For ultimate limit state (ULS) checks of the tower, partial safety factors $\gamma_F = 1.25$ (load factor) and $\gamma_M = 1.1$ (material factor) are adopted in accordance with DNVGL-RP-C202, giving a design yield strength:

$$f_{yd} = \frac{f_y}{\gamma_F \gamma_M},$$

Here, f_y is the material characteristic yield strength. Finally, the utilisation check is expressed as

$$U_{\text{yield}} = \frac{\sigma_{\text{VM,Sd}}}{f_{yd}} \leq 1.0.$$

Appendix E

Verification numerical results for models C1 and C2

This appendix presents numerical results obtained during the dynamic model verification process, as introduced in Section 6.3.1, in terms of the mean responses and loads (Table E.1) and the standard deviation of the responses and loads (Table E.2)

Table E.1: Verification results in operational cases for the baseline FD model (C1) against OpenFAST (OF). Mean of motions and loads.

Case	Surge (m)		Heave (m)		Pitch (deg)		T2 (kN)		My (MNm)		Fxaero (kN)	
	OF	C	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1
3.1	15.54	15.98	0.65	0.66	2.86	2.91	1093.16	1090.81	49.9	52.3	401.14	427.55
3.2	26.84	27.24	1.00	1.01	4.96	4.87	1255.06	1248.68	86.3	90.5	679.69	718.08
3.3	15.27	15.98	0.63	0.66	2.82	2.91	1099.87	1079.01	49.2	52.8	364.64	417.80

Appendix E. Verification numerical results for models C1 and C2

Table E.2: Verification results in wave-only, wind-only, and operational cases for the baseline FD model (C1) against OpenFAST (OF). Standard deviation of motions and loads.

Case	Surge (m)		Heave (m)		Pitch (deg)		Nacelle Accel. (m s^{-2})		T2 (kN)		My (kNm)	
	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1	OF	C1
1.1	0.08	0.07	0.01	0.01	0.04	0.04	0.0936	0.0910	2.56	0.49	3.52	3.57
1.2	0.15	0.15	0.02	0.02	0.08	0.08	0.1872	0.1830	4.58	0.98	7.03	7.13
1.3	0.22	0.21	0.05	0.05	0.11	0.10	0.0786	0.0748	1.28	1.77	3.36	3.33
1.4	0.45	0.42	0.09	0.10	0.22	0.20	0.1572	0.1495	2.49	3.54	6.73	6.66
1.5	0.55	0.56	0.29	0.31	0.42	0.45	0.0666	0.0650	3.24	10.10	4.76	5.36
2.1	0.87	0.84	0.01	0.01	0.13	0.12	0.0050	0.0045	9.48	7.52	2.22	2.31
2.2	2.71	2.90	0.07	0.06	0.41	0.39	0.0216	0.0211	39.15	44.49	6.98	7.54
2.3	3.94	4.29	0.16	0.14	0.68	0.63	0.0386	0.0384	70.48	140.65	11.0	12.3
2.4	4.50	4.84	0.13	0.12	0.81	0.76	0.0558	0.0522	66.55	77.32	12.9	14.5
3.1	2.69	2.74	0.07	0.06	0.42	0.39	0.1374	0.1327	39.07	41.29	8.62	9.16
3.2	3.91	4.06	0.16	0.13	0.68	0.63	0.1766	0.1734	70.25	124.07	12.8	13.9
3.3	4.46	4.48	0.13	0.11	0.81	0.75	0.2574	0.2494	66.23	66.45	16.0	17.4

Table E.3: Verification results in operational cases for the FD+TD model (C2) against OpenFAST (OF). Mean of loads.

Case	T2 (kN)		My (MNm)	
	OF	C2	OF	C2
3.1	1093.16	1092.44	49.9	56.0
3.2	1255.06	1253.43	86.3	94.2
3.3	1099.87	1083.84	49.2	56.6

Table E.4: Verification results in wave-only, wind-only, and operational cases for the FD+TD hybrid model (C2). Standard deviation of loads.

Case	T2 (kN)		My (MNm)	
	OF	C2	OF	C2
1.1	2.56	0.32	3.52	3.54
1.2	4.58	0.62	7.03	7.08
1.3	1.28	1.19	3.36	3.31
1.4	2.49	2.39	6.73	6.63
1.5	3.24	4.56	4.76	5.34
2.1	9.48	9.80	2.22	2.26
2.2	39.15	45.54	6.98	7.20
2.3	70.48	85.33	11.0	11.5
2.4	66.55	76.85	12.9	13.4
3.1	39.07	42.47	8.62	8.80
3.2	70.25	79.99	12.8	13.2
3.3	66.23	70.03	16.0	16.3

Appendix F

Initial mooring surrogate dataset

This appendix presents the initial dataset used to train the initial mooring surrogate model. The full factorial design of experiments over all input parameters listed in Table F.1 yielded 117 187 500 combinations.

Table F.1: Primary platform, mooring system and environmental parameters (design of experiments).

	Min	Max	Units	Count
Platform and mooring design variables				
Nominal chain diameter (d_{moor})	0.04	0.2	m	5
Fairlead radius (r_F)	3.0	10.0	m	5
Fairlead z -coordinate (z_F)	-100.0	-20.0	m	5
Anchor radius (r_A)	3.0	850.0	m	10
Anchor z -coordinate (z_A)	-320.0	-180.0	m	5
Line length (l_{moor})	100.0	1100.0	m	10
Platform vertical centre of mass (z_{CM})	-100.0	-20.0	m	5
Environmental conditions variables				
Mean wind speed (U)	3.0	25.0	m s^{-1}	5
Turbulence intensity (TI)	12	16	-	3
Significant wave height (H_s)	1.0	6.0	m	5
Peak period (T_p)	7.0	10.0	s	5

Given that only about 300 000 samples could be simulated within the available time

Appendix F. Initial mooring surrogate dataset

frame, this was insufficient to achieve a surrogate that would be precise enough. The lack of precision was evident in the verification results, as illustrated in Figure F.1.




Figure F.1: Initial mooring surrogate: true versus predicted value plots and absolute relative error histograms.

Appendix G

ICR time series in regular waves

Figures G.1 and G.2 present the time histories of x_{ICR} and z_{ICR} for two wave periods (10s and 125s), and for four wave heights (2m, 4m, 4m). Significant differences in the sensitivity of the ICR to the wave height is seen between the two periods' cases. The vast discrepancies at low-frequency (right) are not seen in the ICR history in the high-frequency case (left).




Figure G.1: x-coordinate of ICR as excited by high-frequency (left) and low-frequency (right) waves of three different heights.

Appendix G. ICR time series in regular waves




Figure G.2: z -coordinate of ICR as excited by high-frequency (left) and low-frequency (right) waves of three different heights.

Appendix H

ICR in 1D and 2D coupled motion

Figure H.1 compares pitch motion time history in a 2 degrees of freedom (DOF) coupled motion (surge and pitch) against a 1 DOF (pitch only) uncoupled motion. Without the coupling, the pitch motion is not distorted and follows a sinusoidal pattern.




Figure H.1: A comparison of pitch motion time history and FFT amplitudes in a 2 DOF coupled motion against a 1 DOF uncoupled motion.

Appendix I

Sensitivity of ICR to environmental loading parameters

This appendix presents the figures supporting the discussion in Section 5.4. Figure I.1 shows the ICR coordinate distributions for two additional wave periods: 31 s and 125 s (pitch and surge natural frequencies). Figure I.2 shows the distributions of the horizontal velocity at points A and B for all wave periods.

Appendix I. Sensitivity of ICR to environmental loading parameters




Figure I.1: Distribution of ICR coordinates in regular waves of periods 5 s–31 s (top) and 5 s–125 s (bottom).




Figure I.2: Distributions of the horizontal velocity at points A (left) and B (right).

Appendix J

Correlation between design variables and FOWT responses

This appendix presents the correlation matrices (Spearman correlation) between the design variables and the responses of interest for the spar (Fig. J.1) and semisubmersible (Fig. J.2) floating offshore wind turbines.




Figure J.1: Correlation matrix between the design variables and the responses of interest for the spar design.

Appendix J. Correlation between design variables and FOWT responses




Figure J.2: Correlation matrix between the design variables and the responses of interest for the semisubmersible design.

Appendix K

Example of an IPOPT solution

This appendix provides a simple analytical example to illustrate the basic principles of a primal–dual interior-point method, as employed by IPOPT.

K.0.1 Problem definition

Consider the following nonlinear optimisation problem with a single inequality constraint:

$$\min_x \quad f(x) = (x - 2)^2, \quad (\text{K.1})$$

$$\text{s.t.} \quad x \geq 1. \quad (\text{K.2})$$

The unconstrained minimum of the objective function occurs at $x = 2$, which also satisfies the constraint. Therefore, the constrained optimum is $x^* = 2$, and the inequality constraint is inactive at the solution.

K.0.2 Slack-variable formulation

The inequality constraint is converted into an equality constraint by introducing a slack variable $s > 0$:

$$x - 1 - s = 0, \quad s > 0. \quad (\text{K.3})$$

Appendix K. Example of an IPOPT solution

The optimisation variables now consist of the primal variables (x, s) , while the equality constraint is enforced through a Lagrangian formulation.

K.0.3 Barrier problem and Lagrangian formulation

Using an interior-point approach, the inequality constraint is enforced by augmenting the objective function with a logarithmic barrier term. The resulting barrier subproblem is defined as

$$\min_{x,s} \quad \phi_\mu(x, s) = (x - 2)^2 - \mu \ln(s), \quad (\text{K.4})$$

subject to

$$x - 1 - s = 0, \quad (\text{K.5})$$

where $\mu > 0$ is the barrier parameter. To account for the equality constraint, a Lagrangian function is introduced:

$$\mathcal{L}(x, s, \lambda) = (x - 2)^2 - \mu \ln(s) + \lambda(x - 1 - s), \quad (\text{K.6})$$

where λ is the Lagrange multiplier associated with the equality constraint.

K.0.4 KKT conditions

The corresponding Karush–Kuhn–Tucker (KKT) conditions are

$$\frac{\partial \mathcal{L}}{\partial x} = 2(x - 2) + \lambda = 0, \quad (\text{K.7})$$

$$\frac{\partial \mathcal{L}}{\partial s} = -\frac{\mu}{s} - \lambda = 0, \quad (\text{K.8})$$

$$x - 1 - s = 0, \quad (\text{K.9})$$

$$s > 0. \quad (\text{K.10})$$

Combining the first two equations yields the perturbed complementarity condition

$$s\lambda = \mu. \quad (\text{K.11})$$

Appendix K. Example of an IPOPT solution

Together, these equations constitute the primal–dual system solved by the interior-point method.

K.0.5 Central path and convergence

Eliminating λ and using the feasibility condition $s = x - 1$, the KKT system reduces to

$$2(x - 2)(x - 1) = \mu. \quad (\text{K.12})$$

This equation defines the so-called *central path*, along which the iterates (x, s, λ) evolve as the barrier parameter μ is reduced. As $\mu \rightarrow 0$, the solution of the barrier problem converges smoothly to the constrained optimum $x^* = 2$ while maintaining strict feasibility ($s > 0$) at all intermediate iterations. To illustrate this convergence, Table K.1 lists several points on the central path for decreasing values of the barrier parameter μ . For each μ , the corresponding primal and dual variables are obtained from $2(x - 2)(x - 1) = \mu$, $s = x - 1$, and $\lambda = -\mu/s$.

Table K.1: Example points on the central path for decreasing barrier parameter μ .

Iteration	μ	x	s	λ
1	1	2.3660	1.3660	−0.7321
2	10^{-1}	2.0477	1.0477	−0.09545
3	10^{-2}	2.0050	1.0050	−0.009950
4	10^{-3}	2.0005	1.0005	−0.0009995

At the early stages of the optimisation, the barrier parameter is relatively large, strongly penalising small values of the slack variables, discouraging the iterates from approaching the constraint boundaries too rapidly. In this phase, the optimisation primarily seeks directions that reduce the objective function, leading to noticeable changes in the design variables. As the optimisation progresses, IPOPT gradually reduces the barrier parameter. This relaxation diminishes the influence of the barrier term in the merit function, allowing the slack variables to decrease and the solution to approach the active constraint boundaries. Consequently, constraint satisfaction and complementarity conditions become increasingly important in determining the search

Appendix K. Example of an IPOPT solution

direction. In this regime, relatively small changes in the design variables may lead to significant changes in constraint activity, and the optimisation behaviour becomes dominated by feasibility restoration rather than further objective improvement.

In this simple example, the KKT system can be solved analytically for each value of μ , yielding the points along the central path directly. In practice, IPOPT computes these solutions numerically by applying Newton steps to the perturbed KKT system.

K.0.6 Newton solution

The KKT conditions can be written compactly as a nonlinear residual system

$$F(x, s, \lambda) = \begin{bmatrix} 2(x - 2) + \lambda \\ -\mu/s - \lambda \\ x - 1 - s \end{bmatrix} = \mathbf{0}. \quad (\text{K.13})$$

At a current iterate (x^k, s^k, λ^k) , a Newton step $(\Delta x, \Delta s, \Delta \lambda)$ is computed by solving the linear system

$$J_F(x^k, s^k, \lambda^k) \begin{bmatrix} \Delta x \\ \Delta s \\ \Delta \lambda \end{bmatrix} = -F(x^k, s^k, \lambda^k), \quad (\text{K.14})$$

where J_F denotes the Jacobian of the KKT residuals with respect to (x, s, λ) . The iterates are then updated according to

$$(x, s, \lambda)^{k+1} = (x, s, \lambda)^k + \alpha \begin{bmatrix} \Delta x \\ \Delta s \\ \Delta \lambda \end{bmatrix}, \quad (\text{K.15})$$

where the step length $\alpha \in (0, 1]$ is chosen to ensure sufficient decrease of the residual norm while maintaining strict feasibility of the slack variable ($s > 0$).

In practical large-scale problems, this Newton-based procedure is applied repeatedly for a sequence of decreasing barrier parameters μ , yielding convergence to a solution of the original constrained optimisation problem.