Mooring Line Modelling and Design Optimization of Floating Offshore Wind Turbines

by

Matthew Thomas Jair Hall
B.Sc., University of New Brunswick, 2010

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF APPLIED SCIENCE

in the Department of Mechanical Engineering

Mooring Line Modelling and Design Optimization
of Floating Offshore Wind Turbines

by

Matthew Thomas Jair Hall
B.Sc., University of New Brunswick, 2010

Supervisory Committee

Dr. Brad Buckham, Supervisor
(Department of Mechanical Engineering)

Dr. Curran Crawford, Supervisor
(Department of Mechanical Engineering)

Supervisory Committee

Dr. Brad Buckham, Supervisor
(Department of Mechanical Engineering)

Dr. Curran Crawford, Supervisor
(Department of Mechanical Engineering)

ABSTRACT

Floating offshore wind turbines have the potential to become a significant source of affordable renewable energy. However, their strong interactions with both wind- and wave-induced forces raise a number of technical challenges in both modelling and design. This thesis takes aim at some of those challenges.

One of the most uncertain modelling areas is the mooring line dynamics, for which quasi-static models that neglect hydrodynamic forces and mooring line inertia are commonly used. The consequences of using these quasi-static mooring line models as opposed to physically-realistic dynamic mooring line models was studied through a suite of comparison tests performed on three floating turbine designs using test cases incorporating both steady and stochastic wind and wave conditions. To perform this comparison, a dynamic finite-element mooring line model was coupled to the floating wind turbine simulator FAST. The results of the comparison study indicate the need for higher-fidelity dynamic mooring models for all but the most stable support structure configurations.

Industry consensus on an optimal floating wind turbine configuration is inhibited by the complex support structure design problem; it is difficult to parameterize the full range of design options and intuitive tools for navigating the design space are lacking. The notion of an alternative, "hydrodynamics-based" optimization approach, which would abstract details of the platform geometry and deal instead with hydrodynamic performance coefficients, was proposed as a way to obtain a more extensive and intuitive exploration of the design space. A basis function approach, which represents the design space by linearly combining the hydrodynamic performance coefficients

of a diverse set of basis platform geometries, was developed as the most straightforward means to that end. Candidate designs were evaluated in the frequency domain using linearized coefficients for the wind turbine, platform, and mooring system dynamics, with the platform hydrodynamic coefficients calculated according to linear hydrodynamic theory. Results obtained for two mooring systems demonstrate that the approach captures the basic nature of the design space, but further investigation revealed limitations on the physical interpretability of linearly-combined basis platform coefficients..

A different approach was then taken for exploring the design space: a genetic algorithm-based optimization framework. Using a nine-variable support structure parameterization, this framework is able to span a greater extent of the design space than previous approaches in the literature. With a frequency-domain dynamics model that includes linearized viscous drag forces on the structure and linearized mooring forces, it provides a good treatment of the important physical considerations while still being computationally efficient. The genetic algorithm optimization approach provides a unique ability to visualize the design space. Application of the framework to a hypothetical scenario demonstrates the framework's effectiveness and identifies multiple local optima in the design space – some of conventional configurations and others more unusual. By optimizing to minimize both support structure cost and root-mean-square nacelle acceleration, and plotting the design exploration in terms of these quantities, a Pareto front can be seen. Clear trends are visible in the designs as one moves along the front: designs with three outer cylinders are best below a cost of \$6M, designs with six outer cylinders are best above a cost of \$6M, and heave plate size increases with support structure cost. The complexity and unconventional configuration of the Pareto optimal designs may indicate a need for improvement in the framework's cost model.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

I'm grateful to all the people who helped make my time at UVic a rich learning experience. Thanks to my supervisors for providing the chance to work in such an interesting research area. Special thanks to Brad for his fantastic explanations of the range of technical topics that came up. Special thanks to Curran for his attentiveness and ability to provide ideas and direction with perspective and clarity. Thanks to the colleagues-cum-friends I've met at various conferences, especially the INORE symposium, for the good discussions, fun times, and inspiration of knowing how many fantastic people are working in similar research directions. Thanks to Sebastien Gueydon of MARIN for his expert advice which, delivered over the course of just two conferences, helped immensely in keeping my research in touch with reality. Thanks to Scott Beatty for being my floating structure hydrodynamics comrade at UVic and for the opportunity to participate in WEC testing in St. John's. Thanks to all my friends at UVic and elsewhere, for making school and everything else fun. Thanks to my mom and dad for the blessing of their never-failing support. And thanks to my brother, Stef, for his courage and never-failing positivity.

DEDICATION

To the idea that all human construction rests on a foundation of natural ecosystems, and that these ecosystems should – and must – be sustained.

Nomenclature

Greek

β	incident wave heading
ζ	wave elevation
ζ_d	damping ratio of platform motions
ξ	body displacement vector
$\mathbf{\Xi}$	complex amplitude vector of body displacement
ρ	water density
$\sigma_{a\,nac.}$	standard deviation of fore-aft nacelle acceleration
$\bar{\xi_5}$	static or mean pitch angle
ϕ_I	incident wave velocity potential
ϕ_D	wave diffraction velocity potential
ϕ_R	wave radiation velocity potential
ω	wave or platform motion frequency

Latin

$\mathbf{A}$	hydrodynamic added mass matrix
A_{wave}	incident wave amplitude
$\mathbf{B}$	system damping matrix
$\mathbf{B}_{visc.}$	linearized viscous hydrodynamic damping matrix
$\mathbf{C}$	system restoring stiffness matrix
C_A	empirical added mass coefficient
C_D	empirical damping coefficient
D	cylinder diameter
f	hydrodynamic force calculated using the approach of Morison's equation
$\mathbf{f}_e$	excitation force vector on platform from incident waves
$\mathbf{F}_e$	complex amplitude vector of forcing on platform from incident waves
F_S	factor of safety
$\mathbf{f}_r$	forcing on platform from wave radiation
H_I	draft of central cylinder
H_O	draft of outer cylinders
H_s	significant wave height

I_i	system mass moment of inertia in DOF i
I_{wp}	platform water plane moment of inertia in pitch direction
J	objective function value
K_C	Keuleghan-Carpenter number
$\mathbf{k}_e$	wave excitation force kernel (impulse response function) vector
$\mathbf{K}_r$	radiation kernel (impulse response function) matrix
$L_{unstr.}$	unstretched mooring line length
$\mathbf{M}$	system mass matrix
N_F	number of outer cylinders in platform
R_F	horizontal distance from platform center to outer cylinder centers
R_I	radius of central cylinder
R_O	radius of outer cylinders
RAO	response amplitude operator
$S(\omega)$	wave power spectral density function
T	wave or platform motion period
T_I	taper ratio of central cylinder
T_p	peak spectral period
T_{line}	mooring line tension measured at fairlead
u	water velocity
$\forall$	displaced volume of platform
v	body velocity
x	longitudinal axis in the inertial coordinate system
x_M	mooring line configuration decision variable
x_i	basis platform decision variables
$\mathbf{X}(\omega)$	platform wave excitation coefficients vector
y	lateral axis in the inertial coordinate system
W	weighting for multi-objective optimization
w	weighting for metocean conditions
z	vertical (positive-up) axis in the inertial coordinate system

Acronyms

DOF	degree of freedom
OC3	Offshore Code Comparison Collaboration

PSD	power spectral density
RAO	response amplitude operator
RMS	root mean square
TLP	tension leg platform

Chapter 1

Introduction

1.1 Background

Public awareness of the serious climate changes resulting from industrialized nations' voracious use of fossil fuels is dawning. With that awareness comes a motivation to mitigate further climate change. But, with the planet's population soaring past 7 billion people, many regions experiencing rapid improvements in living standards with concomitant increases in per-capita energy consumption, and the depletion of conventional oil reserves causing a shift to unconventional reserves with lower energy returns on investment, the imperative of *reducing* global greenhouse gas emissions is a monumental challenge.

The most direct way to reduce greenhouse gas emissions is to address the source of the problem: reduce the use of fossil fuels. Doing so is also seen to have benefits in terms of energy security, by reducing dependence on imported oil. Given the heavy reliance on cheap abundant energy in developed countries and the significant political, as well as ethical, obstacles to reducing energy use in developing countries, a lot of hope has been placed on renewable energy technologies as a way to displace fossil-based energy sources without curtailing overall energy use.

Among renewable energy technologies, wind energy conversion is one of the most cost-effective and well-established, with an affordability and global installed capacity second only to large-scale hydro [1, 2]. The wind energy industry has reached a level of maturity such that the converter technology has converged to the familiar three-blade horizontal-axis configuration and wind farms, both on land and in shallow water offshore, are built at industrial (MW) rather than experimental (kW) scales.

However, global wind energy installed capacity is still only a fraction of what it could be. This is partly because of:

1. the competitive economics of existing generation sources,
2. concerns about electrical grid operation with increased penetration levels of intermittent wind power,
3. difficulties in siting wind farms due to wind resource limitations, and
4. public resistance to wind turbines near populated areas or in pristine environments.

To combat these factors, the wind energy industry began expanding into the offshore domain in the last two decades, seeking the greater location availability and stronger winds found over water. Offshore wind sites can often be closer to major coastal population centers than onshore wind sites [3]. In addition, by maintaining a minimum distance from shore, can avoid noise and aesthetic constraints, allowing for higher-performance designs and larger-scale installations. Larger device sizes are further enabled by the relative ease with which large structures can be transported over water compared to over land. Finally, since offshore wind speeds are generally higher and turbulence levels lower [4], greater capacity factors can be realized. All these factors combine to allow for more efficient and cost-effective wind turbines, provided challenges associated with the harsh operating environment and offshore construction and maintenance can be overcome [3].

The greatest limitation with conventional offshore wind turbines pertains to water depth; conventional monopile foundations are limited to water depths of about 30 m. More advanced tripod or jacket-type structures are limited to 60 m of depth. In greater depths, structural requirements make the designs economically infeasible [5]. This has greatly limited the extent of the offshore domain in which wind energy can be harnessed.

However, using a floating turbine support structure rather than a bottom-fixed foundation pushes back the depth limitation to hundreds of meters, drastically increasing the siting options while potentially simplifying installation, maintenance, and decommissioning operations. With mooring lines and power transmission cables being the only connection to the sea floor, floating wind turbines can benefit being assembled at-shore and then towed to site, avoiding the complications of on-site

assembly. Similarly, floating wind turbines could be unmoored and towed to shore for major service operations or decommissioning. For these reasons, floating wind turbines hold a lot of potential for increasing wind power generation capacity.

1.2 The Floating Wind Turbine Design Problem

Floating wind turbines face a unique set of design challenges arising from the combined aerodynamic and hydrodynamic factors involved. From a naval architecture perspective, the presence of the wind turbine gives the otherwise-straightforward floating structure an unprecedented sensitivity to wind loadings. If not for the aerodynamic influence of the wind turbine, the floating structure design problem would be very similar to that of an offshore oil platform. In fact, many of the design principles applied to floating wind turbine hydrodynamics come from the offshore oil and gas industry [6].

From a wind turbine design perspective, the use of a floating platform exposes the wind turbine to significant new motions and loads. The six degrees of freedom (DOFs) enabled by the floating platform are pictured in Figure 1.1. Motions in these DOFs are excited by both wave forces on the platform and wind forces on the turbine. The DOF introduced by the floating platform that is most problematic to a turbine is pitch – a fore-aft tilting of the platform. This is the DOF most easily excited by both wave loadings and wind thrust loadings and it influences the bending moments in the tower and the blades – two of the most critical structural loads. Rather than expend resources creating vastly-stronger turbines to handle the loads from these new motions, designers have focussed on designing the support structure (floating platform and mooring system) to minimize platform motions. This is in fact the dominant challenge in floating wind turbine support structure design.

To meet this challenge, the support structure needs to (1) resist the overturning moment caused by steady and time-varying thrust forces of the wind turbine and (2) resist or avoid the motions caused by waves on the platform. These two objectives often conflict with each other, or jointly compete with cost objectives. For example, a high stiffness in pitch is necessary to resist the overturning moment from a steady thrust on the wind turbine. This is most intuitively achieved by using a wide floating platform with a large water plane moment of inertia, providing a large hydrostatic stiffness in the pitch DOF. As soon as waves are added to the picture, however, this type of platform will exhibit significant wave-induced motions. There are two

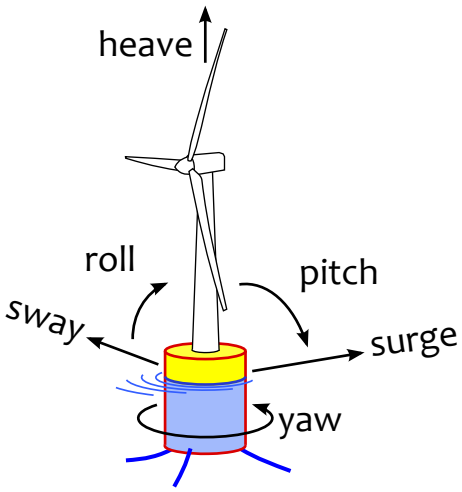

Figure 1.1: Degrees of freedom of a floating wind turbine

alternatives to water plane area for providing the required pitch stiffness: ballast or mooring lines. Ballast can be used with a deep-drafted platform design to lower the center of gravity well below the center of buoyancy in order to provide a large pitch restoring force. Alternatively, taut vertical mooring lines can be used with a submerged, overly-buoyant platform to provide a high stiffness in pitch. Both of these approaches mitigate problems with wave-induced motions by permitting a small water plane area. However, this is achieved at the expense of increasing the displaced volume required of the platform, potentially increasing costs.

1.2.1 Stability Classes

The three means of achieving static stability make up the three "stability class" into which floating wind turbine support structures can be categorized. These classes are illustrated in Figure 1.2.

Buoyancy-stabilized designs rely on a large water plane area to raise the platform's metacenter above its center of gravity. These designs are generally shallow-drafted with a center of mass near the waterline. Common designs are rectangular or circular "barges" or a ring of three or more vertical columns connected together (often called "semi-submersibles"). The shallow drafts of buoyancy-stabilized platforms make for simple installation and the greatest siting flexibility. The lack of ballast reduces size and material requirements. The trade off is that the large water plane area can make the platform more susceptible to wave-induced motions. Heave plates, horizontal

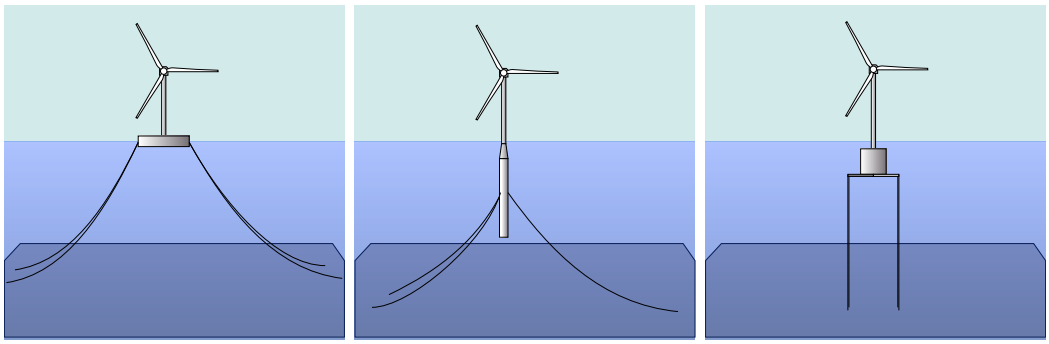

(a) buoyancy-stabilized (b) ballast-stabilized (c) mooring-stabilized

Figure 1.2: Floating wind turbine stability classes

discs for increasing hydrodynamic damping in the vertical direction are often added to the bottom of multi-cylinder semi-submersible platforms to reduce wave-induced motions.

Ballast-stabilized designs rely on a deep draft and heavy ballast to make the platform's center of gravity lie below its center of buoyancy, thus ensuring hydrostatic stability in all circumstances. Because of the draft requirement, these designs almost always use a long vertical cylinder or spar shape, and are called spar-buoys. With a minimal water plane area, a spar-buoy is minimally-susceptible to wave induced motions, but the amount of ballast required adds material and size to the design, raising costs, and the large draft constrains siting and installation options.

Mooring-stabilized platforms, often called tension leg platforms (TLPs), make use of tensioned usually-vertical mooring lines to hold the platform below the waterline, providing a pretension to resist any heaving or pitching motions. With a small water plane area (generally the platform is completely submerged) and highly-tensioned mooring lines, the TLP design is extremely stable. Its disadvantages primarily involve loads and costs associated with the high-tension mooring system, and increased material costs from the extra buoyancy needed to counter the mooring line tension.

1.2.2 Other Considerations

The wind turbine control system also has an important effect on pitch stability. When the turbine is acting at its rated power standard blade-pitch control of the rotor blades can lead to negative damping along the fore-aft axis of the turbine. In this mode, a typical controller will pitch the blades to reduce thrust when it sees an increase in wind speed and increase thrust when it sees a decrease in wind speed, in order to

maintain a constant power output. This change in the thrust force is in the same direction as the hub's motion, resulting in a negative damping effect. Various control schemes have been suggested to avoid this issue, and it is an important research area. Some of the ideas that have been proposed are:

- increasing the delay in the blade-pitch-to-feather controller to greater than the turbines natural pitch period (this is already in practice for fixed foundation turbines because the tower bending natural period is shorter than the reaction time of the blade pitch-to-feather controller [7]),
- using an additional blade-pitch control loop based on nacelle axial acceleration,
- limiting average power output to less than the rated output of the generator so there is headroom to absorb more power during forward pitching,
- using independent blade pitch control causing corrective asymmetric rotor loading [4], and
- using tuned mass dampers [8].

Many other design considerations unrelated to platform dynamics surround the support structure design problem. Installation requirements can have a big effect on technical and economic feasibility. Shallow-drafted and inherently-stable platforms, typical of buoyancy-stabilized designs, have the advantage that they can be assembled at shore and then towed to location so that the only on-site construction required is for mooring and transmission systems. TLP designs that may be hydrostatically unstable or deep-drafted spar buoy designs that cannot float in shallow water require specialized vessels to transport them to location, adding cost and time to the installation. The relative ease or difficulty of installing the mooring lines and anchors is another factor.

Siting is another important consideration. While floating wind turbines can be situated in water depths much greater than the 60 m limit of bottom-fixed turbines, the *minimum* depth for the support structure must also be considered. A spar buoy with a 120 m draft, for example, cannot be installed in water that is only 100 m deep, so this technology leaves a large gap of intermediate water depths where neither a fixed nor a floating base will work. A buoyancy-stabilized design on the other hand may physically fit in water depths as shallow as 50 m, but careful attention has to be paid to the mooring system to ensure that shallow water waves do not cause snap

loads on the mooring lines. TLP designs face a similar issue, especially in shallow water; if the mooring lines becoming slack, catastrophic snap loads could be seen when they regain tension.

There are also more practical, manufacturing- and deployment-level considerations. Costs are an extremely important factor when dealing with the slim profit margins of wind energy (as opposed to other offshore industries like oil and gas) and the material and manufacturing costs of the support structure are a big part of that. The size, materials, and complexity of the platform all relate to those costs, as do the required strength and length of the mooring lines and the required anchor type.

The wide range of options and considerations in a floating wind turbine support structure make for a very large design space and many factors need to be considered when deciding on a support structure configuration. Unfortunately, the main functions of the support structure – resisting the wind loads and avoiding the wave motions – lead to conflicting design requirements, making for a design problem that is difficult to navigate.

1.3 State of the Industry

1.3.1 Prototyped Designs

Floating wind turbines are not a new idea; they were first proposed in 1972 by University of Massachusetts professor William Heronemus [9]. However, published research was sparse until the 1990s, when many different floating wind turbine designs began to appear. Only recently have there been significant prototype developments.

Blue H designed and built a tension leg platform with a central float and 6 outer columns onto which the tension legs attach, shown in Figure 1.3. An 80 kW prototype was tested off the coast of Italy in 2007 [10].

In 2007, the Norwegian energy company StatoilHydro partnered with wind turbine manufacturer Siemens to develop the first MW-scale floating wind turbine. The resulting design, called Hywind and shown in Figure 1.4, is a spar-buoy design that takes advantage of the stability of a slender deep-drafted spar. The small water plane area makes the natural frequencies of the spar well below wave excitation frequencies, and the deep 117 m draft places the majority of the structure at depths where wave velocities are minimal. Together, these minimize wave-induced motions. Using rock and water ballast, sufficient distance between center of gravity and center of buoyancy

Figure 1.3: The Blue H 80 kW mooring-stabilized prototype

is provided to keep the static tilt from turbine thrust loads small [7]. Pitch instabilities from conventional pitch control at rated power are avoided using a specially-designed blade-pitch controller [11]. Three slack catenary mooring lines attached mid-way down the spar provide station keeping. The mooring lines also provide yaw stiffness by means of a "crow foot" where each line attaches to the spar [7]. The design features a 2.3 MW Siemens offshore turbine. It was commissioned in the fall of 2009 off the coast of Norway [11].

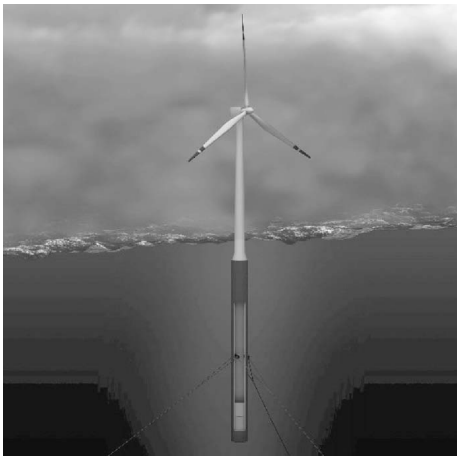

Figure 1.4: The Statiol Hywind 2.3 MW ballast-stabilized prototype

The Danish company Floating Power Plant designed a floating platform called Poseidon that differentiates itself from other floating wind turbine platforms by supporting three wind turbines as well as an array of wave energy converters (WECs). The Poseidon is buoyancy stabilized and is much wider than it is long, in order to support a row of pivoting-float WECs arranged as a terminator arrangement parallel to the wave fronts. The width also provides lateral spacing between the wind turbines. A small-scale Poseidon prototype, shown in Figure 1.5, was deployed in 2008. This prototype is 37 m wide, with three 11 kW two-bladed downwind-rotor wind turbines and 10 WEC units [12].

Figure 1.5: The Floating Power Plant Poseidon 3x11 kW prototype

In 2011, the Norwegian company Sway deployed a scaled prototype of their unique spar-buoy design with a 7 kW wind turbine, shown in Figure 1.6 [13]. The Sway floating wind turbine design makes a number of creative modifications to a more typical spar-buoy design like the Hywind; a downwind coned rotor enables passive yawing and allows the rotor to stay aligned with the wind direction even while the tower is at a large pitch angle under the static thrust load of the turbine. The passive yawing allows for the entire tower to pivot with the turbine, which in turn allows the strategic placement of tension cables along the upwind side of the tower to reduce bending moments in the tower. The lack of required yaw stiffness also enables the use of a single rigid mooring line with a suction-pile anchor. The allowance for high static tilt angles reduces the size and ballast required of the spar.

Also in 2011, the American company Principle Power deployed their first MW-scale prototype, the WindFloat, in Portugal. Originally design by Marine Innovation & Technology, the WindFloat is a buoyancy-stabilized floating wind turbine platform derived from an offshore oil and gas semi-submersible design. Shown in Figure 1.7, it consists of three 11 m-diameter vertical columns submerged to a 23 m draft. To counter wave-induced motions from the large water plane area, heave plates are fitted

Figure 1.6: The SWAY 7 kW prototype

to the bottoms of the cylinders. These plates increase added mass and damping in the vertical direction resulting in longer natural periods and greater damping in the heave, pitch, and roll DOFs [6]. While the support structure is designed for a 5 MW turbine, the prototype uses a 2 MW turbine. The unbalanced weight and variable overturning moments created by changing rotor thrust loads are countered by an active ballast system that pumps water between columns. The WindFloat uses low-tension catenary mooring lines [14].

1.3.2 Conceptual Designs

A number of other noteworthy designs exist that have not yet been realized as prototypes. The Universities of Glasgow and Strathclyde under contract to ITI Energy designed a square buoyancy-stabilized barge platform to support a floating wind turbine and an oscillating water column wave energy device [15]. Though it has been used in a number of modelling studies, the design suffers from large wave-induced motions that make it a poor candidate for further development.

The floating wind turbine research program at MIT has focussed heavily on optimization of a tension leg platform (TLP) design, yielding an MIT TLP design which has been used in a number of modelling studies [16].

All of the above prototypes, and the vast majority of proposed floating wind turbine designs, feature horizontal-axis turbines. There is however some noteworthy research underway on vertical axis turbines specifically for floating offshore applica-

Figure 1.7: The WindFloat buoyancy-stabilized design

tion. One such design is French company Nenuphar Wind's Verti-Wind, pictured in Figure 1.8. Another is the DeepWind concept from Riso-DTU, which uses a Darrieus vertical-axis wind turbine (VAWT) rigidly connected to a spar-buoy float. Unlike typical spar-buoy designs, the moorings are attached at the bottom of the spar, allowing the spar to tilt to absorb the static wind turbine thrust loads, as illustrated in Figure 1.9. A key innovation of the Deep Wind concept is that the generator is located at the bottom of the spar at the level of the mooring attachments, providing much-needed ballast, and the entire spar spins with the Darrieus rotor, eliminating the otherwise-large bending moments at the generator shaft and the need for an expensive main bearing typically found on a VAWT [17], [18].

The wide variety of support platform designs that have been conceived and developed demonstrates the unresolved difficulty in finding an optimal floating wind turbine support structure configuration.

1.3.3 Current Research Areas

With only a handful of kW-scale prototypes and two MW-scale prototypes in the water, the floating wind turbine industry is still in its infancy. While a handful of developers (of various sizes) work on developing and realizing their design concepts (with varying levels of success), the majority of research by academia and interested

Figure 1.8: The Verti-Wind floating VAWT design

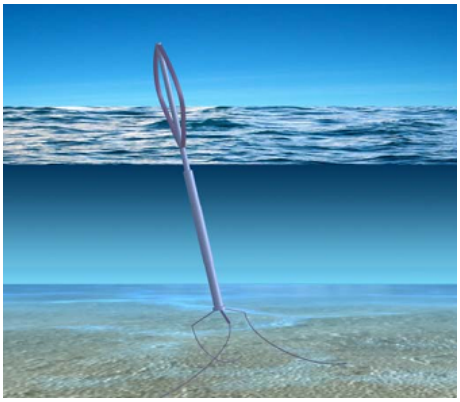

Figure 1.9: The DeepWind floating VAWT design

parties in the larger offshore industry focusses on improving computer modelling capabilities and exploring different design options that have potential to improve the outlook of floating wind.

Given the expense of testing large-scale floating wind prototypes, the few floating wind turbines currently in existence, and the difficulties with smaller-scale model testing arising from scaling incompatibilities between Reynolds number and Froude number, computer modelling is relied on extensively for evaluating floating wind turbine designs. Because the highly-coupled nature of the aerodynamics and hydrodynamics of a floating wind turbine is relatively unique, and conventional wind turbine computer models are not suited to deal with large turbine motions, there is a need to improve computer modelling capabilities to the point where their results can be taken with a high level of confidence. Weaknesses exist in all three modelling areas: the

aerodynamics of the turbine, the hydrodynamics of the platform, and the dynamics of the mooring lines. It is not simply a matter of using the best available models, with high-fidelity computational fluid dynamics and finite element dynamics, because these options tend to be slow to execute and simulation speed is key when iterating over design options. An overview of existing modelling options is given in the next chapter.

Research into different design options – from new platform features to new control schemes or the use of tuned mass dampers – relies primarily on existing modelling capabilities, but also can involve the development of new models to support the new features being analyzed. For example: Cermelli and Roddier developed a specialized hydrodynamics model to explore the effects of heave plates as part of their development of the WindFload design [19], Steward and Lackner have explored the use of active tuned mass dampers for reducing floating wind turbine pitching motions [20], Namik and Stol have explored the use of individual blade pitch control for the same purpose [21], and Vita et al. have explored the options for floating vertical axis wind turbines as part of the development of their DeepWind concept [18].

There has also been work done comparing design possibilities and trying to identify an optimal support structure configuration. One of the most notable comparison efforts, by Robertson and Jonkman, simulated six support structure concepts and compared their performance in a number of environmental conditions [22]. To identify new optimal designs, Sclavounos, Tracy, and Lee conducted a parametric exploration of the design space to search for Pareto-optimal support structure designs. Their work is probably the most notable attempt at global support structure design space exploration to-date, but their cylindrical platform parameterization excluded a significant range of buoyancy-stabilized platform configurations [16]. The floating wind research field has not yet included any truly broad design space exploration studies.

1.4 Key Contributions

The overarching intention of this thesis is to develop computational design tools that will uncover features of the support structure design problem and aid the development of floating wind turbines as a significant source of renewable energy. The work has centered around two distinct focii: evaluating the suitability of alternative mooring line models, and finding globally-optimal support structure configurations from across

the broad design space. With these focii, the work has made two main contributions:

1. **Quantified the effect of mooring line model fidelity on the accuracy of floating wind turbine simulations under various conditions**.
    - Created a high-accuracy time-domain mooring line simulation capability by coupling a high-fidelity mooring line model to an industry standard floating wind turbine simulation code.
    - Simulated a selection of floating wind turbine designs in a range of conditions using both the new high-fidelity mooring line modelling capability and a conventional lower-fidelity mooring line model.
    - Compared the results from the two different mooring line models to determine the effect the choice of mooring model has on the accuracy of the simulation.
2. **Developed an optimization framework that enables a broader exploration of the support structure design space than previously possible**.
    - Devised a platform geometry and mooring system parameterization scheme, with accompanying cost function, that captures a wide extent of the desired design space using a minimal set of decision variables.
    - Synthesized a coupled frequency-domain model that can be used to evaluate the performance of designs described by the parameterization scheme; incorporated wind turbine aerodynamics, platform hydrodynamics, and mooring line dynamics at an appropriate level of fidelity.
    - Developed an optimization algorithm that is adept at navigating the unique decision variable relationships and the discontinuities that exist over the full extent of the floating support structure design space.
    - Demonstrated the operation and utility of the framework through a case study optimization over a realistic site scenario and verified selected results using a higher-fidelity time-domain simulation.

1.5 Thesis Outline

The remainder of this thesis is organized into the following chapters:

Chapter 2 introduces the features of coupled floating wind turbine simulations and describes the range of modelling tools that pertain to floating wind turbine analysis, some of which are central to the work in later chapters.

Chapter 3 presents a study on the importance of mooring model fidelity to floating wind turbine simulation, which embodies the first contribution of this thesis.

Chapter 4 presents a novel attempt at global support structure optimization using basis functions to avoid the limitations of geometry-based decision variables.

Chapter 5 presents a genetic algorithm-based support structure optimization scheme, which embodies the second contribution of this thesis.

Chapter 6 contextualizes the overall conclusions from the research and discusses advisable directions for research continuation.

Chapter 2

Floating Wind Turbine Modelling

2.1 Introduction to Coupled Floating Wind Turbine Simulation

The simulation of floating wind turbines draws on a number of engineering disciplines, principally: aerodynamics, structural dynamics, hydrodynamics, and controls. Models incorporating the relevant factors from each of these disciplines need to be coupled together in order to provide a working model of the entire floating wind turbine system – from the mooring lines to the floating platform to the rotor blades, and all the structural elements and controlled electromechanical systems that connect them. An understanding of each of these disciplines and their associated modelling techniques is critical to the effective development, application, and analysis of floating wind turbine simulations. Figure 2.1 describes the important loads affecting a floating wind turbine structure.

An important distinction among models is whether they are based in the time domain or the frequency domain. Frequency-domain models, which are built on the assumption of periodic harmonic motion, are generally simpler to create, understand, and work with, and their computational simplicity makes them ideal for use "in the loop" of design iteration or optimization task. However, their assumptions render them applicable for only a limited range of situations. Time-domain models, without such limiting assumptions, can potentially be applicable for any situation and in general provide much more detailed (albeit less concise) results. This comes at the expense of computational speed – a result of modelling based on the discretization of differential equations that need to be evaluated at every simulation time step.

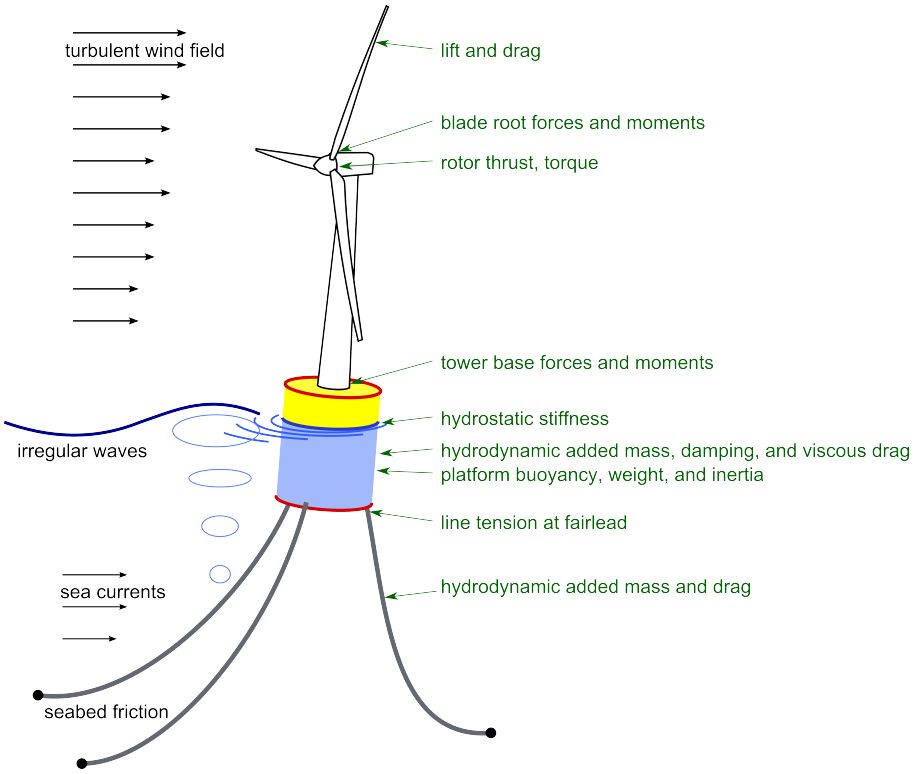

Figure 2.1: Important loads on a floating wind turbine

Some phenomena lend themselves well to frequency-domain modelling while others lend themselves well to time-domain modelling. The hydrodynamics of the floating platform are most conveniently modelled in the frequency domain due to the periodic nature of ocean surface waves. The aerodynamics of the wind turbine, subject to stochastic wind gusts and delayed control adjustments, are most easily modelled in the time domain. In assembling a complete floating wind turbine model, it is often necessary to combine time-domain models and frequency-domain models. This typically entails taking the results from the non-matching model and converting them in a pre-processing step into a form usable by the overall coupled simulation. If the simulation is to be in the frequency domain, a time-domain model needs to be linearized and possibly Fourier transformed – that is, its behaviour analyzed over a range of displacements and periodic motions (if applicable) to obtain linear stiffness, damping, and mass matrices which can then be fed into the frequency-domain equations of motion. If the simulation is to be in the time domain, the output of frequency domain models may need to be put through an inverse Fourier transform and, if any memory effects exist, put through a convolution integral (or alternative treatment)

to produce an impulse response function that can then be applied in the time-domain equations of motion.

Beyond the accuracy of the component models themselves, the accuracy resulting when the models are coupled together is an important issue, especially when frequency- and time-domain conversions are involved. Because inaccuracies arising from one model can affect the dynamics predicted by another model, the significance of a model's limitations needs to be evaluated in the context of a fully coupled simulation rather than in isolation.

The largest organized effort to explore the effects of modelling limitations in the context of floating wind turbine simulation is the the Offshore Code Comparison Collaboration (OC3), an international collaborative effort operating under Subtask 2 of International Energy Agency Wind Task 23 which seeks to compare and verify modelling techniques for offshore wind turbines. The project has gone through four phases comparing, in turn, simulations of a monopile-mounted offshore turbine with a rigid foundation, a monopile-mounted turbine with a flexible foundation subject to soil interactions, a tripod-mounted turbine, and a floating spar-buoy-mounted turbine [23].

Each of these four phases used the NREL offshore 5 MW baseline wind turbine specifications – detailed and publicly-available specifications for a hypothetical wind turbine for use on offshore platforms developed by NREL for the purpose of testing and comparing simulation tools [24]. The turbine design was used extensively in this thesis. Its basic specifications are provided in Table 2.1 and its performance curves are shown in Figure 2.2.

Table 2.1: NREL offshore 5MW baseline wind turbine properties

Property	Value
rated power	5 MW
rotor diameter	126 m
hub height	90 m
cut-in wind speed	3 m/s
rated wind speed	11.4 m/s
cut-out wind speed	25 m/s
rotor mass	110 000 kg
nacelle mass	240 000 kg
tower mass	347 460 kg
center of mass height	64 m

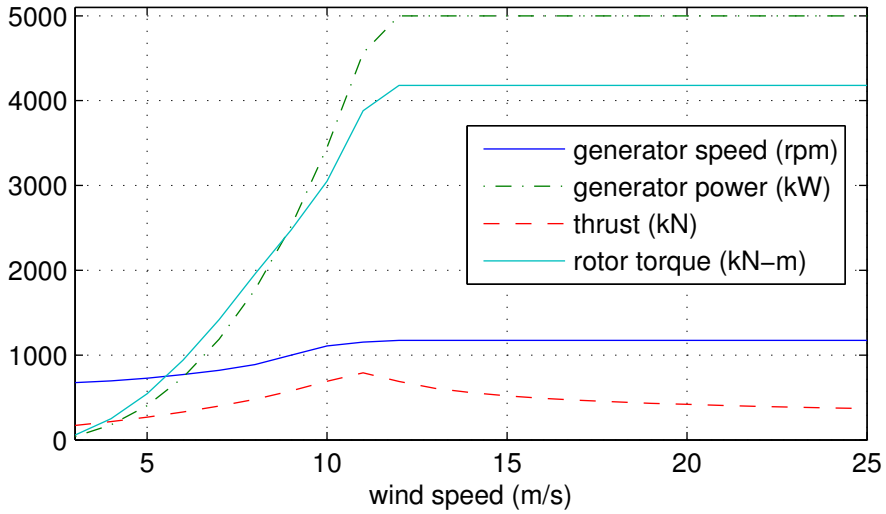

Figure 2.2: Performance curves for the NREL offshore 5MW baseline wind turbine

The emphasis in the OC3 studies is on the offshore support structure dynamics. The OC3 studies make use of a stepwise comparison process – starting with very simple test cases and incrementally adding degrees of freedom and loading sources to distinguish which features cause disagreement between different simulators' results. Environmental inputs (including turbulent wind fields and irregular wave conditions) and turbine and support structure specifications are all standardized between simulators. Phase IV, the phase of interest for floating wind turbine simulation, used a spar-buoy platform configuration based on the Statoil Hywind design. A continuation of the work, called the offshore code comparison collaboration continuation (OC4), is underway to apply the same methods of analysis to a semi-submersible floating platform configuration [25].

The models involved in floating wind turbine simulations can be lumped into three areas: aero-elastic models of the wind turbine and tower, hydrodynamic models of the floating platform, and hydro-elastic models of the mooring system. A range of modelling techniques, with different levels of detail/fidelity and corresponding computational costs, are available for each area.

2.2 Wind Turbine Dynamics Modelling

The wind turbine – comprising the tower, nacelle, and rotor – is a relatively flexible structure and it operates in an unsteady wind field that includes turbulence, a velocity

gradient due to wind shear, and the possibility of airflow approaching at an angle from the rotor axis due to yaw error. Modelling of the wind turbine dynamics generally includes not only an aerodynamic model to account for aerodynamic forces on the blades but also a structural model to account for deflections in the turbine blades and tower, which are far too large to be neglected. The aerodynamic model and structural model are then coupled together, providing feedback to each other during the simulation. Discussion of the wind turbine modelling will be kept brief because the focus in this thesis is on modelling the support structure.

2.2.1 Aerodynamic Models

The common aerodynamic model types are blade element momentum theory (BEM), vortex methods, and computational fluid dynamics (CFD).

Blade Element Momentum Theory

Blade element momentum (BEM) theory is the most common approach to wind turbine aerodynamics modelling because of its superior computational efficiency. It is a coupling of two separate theories: blade element theory and momentum theory. Blade element theory models a blade by dividing it into a number of discrete segments along its length and analyzing the forces on each segment independently using two-dimensional airfoil lift and drag coefficients. It ignores any effect of an element on the flowfield or adjacent elements, and ignores the three-dimensional nature of the blade. Momentum theory is a model for the loss of pressure or momentum across the rotor plane caused by the work done by the air on the rotor. It describes how the rotor alters the velocities in the flowfield. Coupling these two theories together into BEM theory allows the forces on the rotor to be calculated using airfoil lift and drag curves, and the inflow of air on the rotor plane to be adjusted for the effect of the blades. In practice, equations for the two models are solved iteratively. A number of corrections exist to account for hub and tip losses (three-dimensional flow), yawed operation, dynamic stall and large induction factors [26].

Vortex Methods

Vortex methods assume the flow is inviscid and irrotational and use potential flow theory to calculate the airflow and the aerodynamic forces on the blades. Bound vortex filaments situated along the rotor blade span model the lift, and trailing vortex

filaments at the blade root and tip (and possibly at nodes along the blade span) account for the effects of the wake. Velocities induced by these vortex filaments are computed by the Biot-Savart law.

Such vortex methods allow the wake vortex filaments to convect under the influence of the surrounding vorticity, thereby modelling the dynamic nature of the wake. The inherent flexibility of this approach avoids the many correction factors needed in BEM, but adds significant computational complexity [27].

Computational Fluid Dynamics

Computational fluid dynamics (CFD) methods work by discretizing the full Navier-Stokes equations, thereby providing the greatest potential for realistic simulation of the flow field. However, their extreme computational requirements and difficulties associated with numerical issues, modelling flow seperation, and preserving convected vorticity discourage the use of CFD methods for wind turbine aerodynamics [28].

2.2.2 Structural Models

Like the aerodynamic models, existing structural models for the rotor blades and tower span a variety of fidelities. The main structural model types are modal representations, multibody representations, and finite element methods.

Modal Representations

The modal approach represents structural deflections using mode shapes calculated for the flexible members of interest – typically the tower and blades. Knowing the shapes, natural frequencies, stiffness, and damping of the modes of interest is sufficient to create a linear dynamic model of these structural deflections. The minimal computational requirements of these models make them popular in wind turbine simulation, though the assumption of linearity can be questionable for the high deflections and motions possible with floating wind turbines [29].

Multibody Representations

The multibody approach represents flexible structural elements by a series of lumped masses connected by multi-dimensional spring-damper elements. By discretizing a structure in this way, its nonlinear response can be modelled [30]. Because it captures

nonlinear structural dynamics in a flexible and computationally-efficient formulation, this approach is used in a number of wind turbine simulation codes.

Finite Element Methods

The finite element approach, which discretizes flexible structural elements into a mesh of finite structural elements over which the governing equations are applied, gives a very high level of accuracy that is essential for the structural design of wind turbine components. However, its high computational requirements make it inconvenient for coupled floating wind turbine simulation [29].

2.2.3 Current Trends

For aerodynamics modelling, the capacity of BEM theory, the computationally-efficient standard for bottom-fixed wind turbines, is challenged by the large rotor motions that can occur with a floating wind turbine. This, along with other limitations of BEM theory, is motivating the continued development of vortex method-based models. However, BEM theory is still the current standard for floating offshore turbines because of its speed and established reputation. If the motions of the floating support structure can be kept to low levels (as is the goal in the design of such platforms), then the continued use of BEM theory in floating wind turbine simulations may well be adequate.

For structural modelling of the rotor, modal and multibody representations are very popular, and it is rare for the higher fidelity of computationally-expensive finite-element models to be needed for coupled simulations in the course of floating wind turbine design work.

2.3 Platform Hydrodynamics Modelling

The hydrodynamics of the floating platform is a complex fluid-structure interaction problem that includes the effects of excitation from incident ocean waves, damping from waves radiated by the platform, and drag and added mass forces arising from the platform's motion through the water. Structural deflections in the platform are generally negligible compared to the gross motions of the platform. The conventions used for describing the platform motions and the global coordinate system are shown in Figure 2.3. Before discussing the applicable modelling techniques, it is useful

to review how the different hydrodynamic influences on the platform motion are commonly classified.

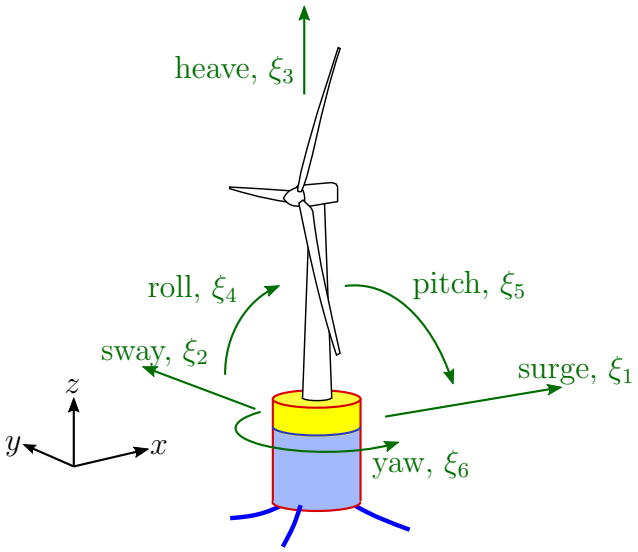

Figure 2.3: Floating wind turbine coordinate system

2.3.1 Hydrodynamic Loadings

The most dominant hydrodynamic loadings on a floating platform come from buoyancy, waves, and the platform's own motion in the water.

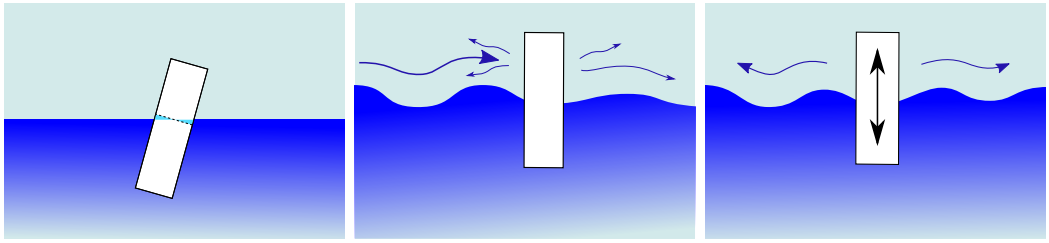

(a) hydrostatics (b) diffraction (c) radiation

Figure 2.4: The components of linear hydrodynamics illustrated for a vertical cylinder.

Hydrostatics (Figure 2.4a) refers to the static restoring force occurring as a result of buoyancy when the platform is displaced along one of its DOFs from its equilibrium position/orientation. As such, it can be calculated by finding the magnitude and centroid of the water volume displaced by the platform.

Wave excitation (Figure 2.4b) is the loading on the platform exerted by incident waves, often without taking into account the motion of the platform. It is commonly

called the diffraction problem because it deals with the force caused by the waves as they diffract around the platform. For platforms that are very small relative to the wavelength, the waves scattering or diffracting around it are minimal and the wave excitation forces can be calculated based on the undisturbed wave kinematics alone.

As a platform moves in the water it generates waves that are radiated outward (Figure 2.4c). Determining the loadings on a platform that result from its motion in the water is referred to as the radiation problem. In an inviscid approximation, the forces associated with wave radiation are distinguished as either added mass or wave-radiation damping. Added mass is the force component in phase with and proportional to the platform's acceleration. It is a result of the mass of water that is accelerated with the platform as the platform moves, and it is frequency dependent. Because of the high density of water, the added mass term can be of the same order as the mass of the platform. (There can be coupling between DOFs where acceleration in one DOF causes a force in another DOF [4].) Wave-radiation damping is the force in phase with and proportional to the platform's velocity (which makes it a linear damping force). In a linear hydrodynamics approach, which ignores viscous drag, the power lost to this damping force is equal to the power being radiated away from the platform in the outgoing waves [31]. The energy in the outgoing waves (and their speed) is dependent on the platform oscillation frequency, and therefore the magnitude of the damping force is frequency dependent. Because the waves generated by platform motion continue even after the platform motion has ceased, wave radiation exhibits a so-called memory effect; instantaneous wave-radiation damping forces depend on past platform motions.

Any relative motion of body and water will also produce viscous drag forces. The relative importance of viscous versus inviscid forces on a submerged body undergoing harmonic motion is represented by the Keulegan-Carpenter number, defined as:

$$K_C = \frac{VT}{D} \tag{2.1}$$

where V is amplitude of oscillation velocity, T is oscillation period, and D is diameter or other characteristic body length. For large K_C, viscous effects dominate and the drag can be approximated using strip theory and Morison's equation (described in Section 2.3.2). For small K_C, inertial forces dominate and the added mass and damping calculated based on wave radiation provide a good approximation. In between, or for complex structures with both large and small bodies, both viscous and inviscid

effects need to be modelled.

Sea currents are another source of loading on the body, in the form of a steady lateral drag force on the platform. The loading from a constant velocity flow of this type is usually treated as a quadratic viscous drag term with an appropriate drag coefficient, using Morison's equation. There are also a number of higher-order hydrodynamic effects resulting from waves on the free surface. Several of the more prominent ones are discussed in Section 2.3.5.

The following subsections discuss the hydrodynamic methods most applicable to floating wind turbine simulation.

2.3.2 Strip Theory and Morison's Equation

Morison's equation is an approach for calculating the transverse hydrodynamic forces on slender cylindrical bodies. It has a nonlinear viscous drag term as well as an added mass term, making it applicable for both steady current-type loads as well as unsteady forces from waves or body motion. The relative form of the equation, accounting for both wave velocity, u, and body velocity, v, is:

$$f = \underbrace{\rho\frac{\pi}{4}D^2H\dot{u}}_{inertia} + \underbrace{\rho\frac{\pi}{4}D^2HC_A(\dot{u}-\dot{v})}_{added\,mass} + \underbrace{\frac{1}{2}\rho DHC_D(u-v)|u-v|}_{drag} \tag{2.2}$$

where f is hydrodynamic force, D is cylinder diameter, H is cylinder draft, and C_A and C_D are empirical added mass and drag coefficients, respectively.

To make it work with rotating degrees of freedom and non-uniform water kinematics, Morison's equation is often combined with strip theory. Strip theory divides the draft of the cylinder into a number of cylindrical strips and analyzes the forces at each strip, an approach analogous to blade-element theory as applied to wind turbine aerodynamics.

$$df = \left(\rho\frac{\pi}{4}D^2\dot{u} + \rho\frac{\pi}{4}D^2C_A(\dot{u}-\dot{v}) + \frac{1}{2}\rho DC_D(u-v)|u-v|\right)dH \tag{2.3}$$

Morison's equation is convenient in that it provides drag and added mass forcings directly from relative fluid velocities and accelerations. However, it is only strictly applicable to slender axisymmetric bodies. This is because it cannot handle coupling between different DOFs in added mass, and it assumes that the body will have negligible effect on the wave motion [29]. This latter assumption is formalized in

G.I. Taylor's long wavelength approximation, which states that for body diameters less than one-fifth the wavelength, diffraction of the wave around the body will be minimal and loadings on the body can be approximated based on the undisturbed wave kinematics [32]. The idea is that if a body is significantly smaller than the wavelength, it will be "transparent" to the wave and have a negligible impact on the wave motion.

A further limitation of Morison's equation is that it includes only a viscous drag (proportional to the square of the velocity) term, thereby assuming that pressure drag from wave radiation is negligible. This only holds true if the movements of the body are very small relative to the body size [4]. When damping from wave radiation is no longer negligible, it can be approximated using the Haskind relation, which relates the excitation force to the radiation damping [33].

Morison's equation can also be adapted to include drag and added mass coefficients for flow along the axis of the cylinder; the drag coefficient can be applied to the surface area of the cylinder and the added mass can be calculated based on the volume of an imaginary half dome formed at the end of the cylinder. Lift forces from asymmetrical flow separation over the cylinder (Karman vortex street) can be calculated using a lift coefficient in Morison's equation and Strouhal number for its frequency of oscillation [33].

2.3.3 Introduction to Linear Hydrodynamics

When the body is large enough relative to the wavelengths that wave diffraction effects become significant, or its movements are large enough that wave radiation effects become significant, Morison's equation is no longer adequate and a theory that includes the effects of the body on the water needs to be used. The simplest such theory is linear hydrodynamics.

Linear hydrodynamics simplifies the calculation of the wave diffraction and wave radiation forces by analyzing them independently. Because the problem is approximated as linear, superposition can be used to break the hydrodynamics forces on the body down into three independent components: hydrostatics (buoyancy), diffraction (wave excitation) and radiation (added mass and wave-radiation damping).

Linear hydrodynamics uses Airy wave theory, which models the hydrodynamics using potential flow. Therefore, (nonlinear) viscous drag is neglected. Furthermore, since the three hydrodynamic components are analyzed independently, it is assumed

that the motion of the body does not affect the wave loading on the body; the displacements of the body must be small relative to the body size and the pitch and roll angles must be small. The standard assumptions in linear hydrodynamics are [4]:

- waves are not steep ($\frac{d\zeta}{dx} \ll 1$, where ζ is free surface elevation)
- flow is inviscid (small K_C)
- body displacements are small relative to wave amplitude
- body displacements are small relative to body size

Hydrostatics

The small-angle and linearity assumptions of linear hydrodynamics theory make for an extremely simple treatment of hydrostatics; the reaction forces/moments for displacement/rotation in each DOF are given by stiffness constants calculated from the water plane geometry. Because of the possibility of coupling between DOFs, the result is a six-by-six hydrostatic stiffness matrix, $\mathbf{C}$, with element i, j representing the force/moment in DOF j from a unit translation/rotation in DOF i. There is no hydrostatic stiffness in the surge, sway, or yaw DOFs as the water plane is invariant in these DOFs.

Wave Diffraction

The wave excitation loading in linear hydrodynamics usually requires a numerical approach for all but the simplest of objects and relies on potential flow theory. In this theory, the water velocity field can be represented by a velocity potential, ϕ, such that $\mathbf{u} = \nabla\phi$.

The surface elevation of a monochromatic linear progressive wave (a travelling sinusoidal wave) is described using Airy wave theory by $\zeta = A_{wave}\cos(\omega t - kx)$, where A_{wave} is wave amplitude, ω is wave frequency, and k is wave number. The complex velocity potential (in deep water) that will produce this wave is

$$\phi_I = Re\left\{\frac{igA_{wave}}{\omega}e^{kz-ikx-i\omega t}\right\}, \tag{2.4}$$

where $i = \sqrt{-1}$ [34].

With the assumptions of linear hydrodynamics, the excitation force from an incident wave can be studied in isolation from the body's movement. If the body is held

still, a "diffraction potential", ϕ_D, is required to make the wave diffract around the body rather than travelling through it. The diffraction potential is the solution to a boundary value problem for the Laplace equation, $\nabla^2(\phi_I + \phi_D) = 0$, with boundary conditions of no flow through the surface of the body:

$$\frac{\partial \phi_I}{\partial \hat{n}} + \frac{\partial \phi_D}{\partial \hat{n}} = 0 \quad \text{on} \quad S_{body} \tag{2.5}$$

where $\hat{n}$ is the vector normal to the body surface, S_{body} [34].

The standard solution approach for such problems is a panel method, in which the surface of the submerged body is discretized into panels, each having unknown potential source strength. The summed solution of each of these panels' source strengths forms the diffraction potential [35].

If the body is relatively slender, with a small diameter relative to the wave lengths encountered, the wave-scattering effect can be neglected, as was discussed in Section 2.3.2. In this case, only the incident wave potential needs to be taken into account, and the scattering potential can be neglected. The resulting simplified excitation force is called the Froude-Krylov force, and is advantageous compared to the full diffraction problem because numerical solution of the scattering potential is no longer required.

In any case, because of the periodic nature of the problem, it can be solved in the frequency domain. Using complex numbers to indicate both magnitude and phase, one boundary value problem fully describes the flow for a given wave frequency and heading.

Once the velocity potential is known around the object, the linear hydrodynamic pressure $p = -\rho\frac{\partial \phi}{\partial t}$ (taken from the Bournoulli equation) can be integrated over the surface of the body to find the resulting wave excitation forces and moments [33]. In practice, these forces and moments are often normalized by wave amplitude to yield frequency-domain wave excitation coefficients for each DOF. The vector of these coefficients is denoted by $\mathbf{X}$.

Wave Radiation

The radiation problem provides the third and final piece of the linear hydrodynamics puzzle. Whereas the diffraction problem analyzes forces from waves while the body is held still, the radiation problem analyzes forces generated when the body is moved in still water.

Waves generated by the body's motion propagate outward over the water surface

even after the body's motion has stopped. These radiated waves influence the pressure field in their wake all the way back to the body, and therefore continue to impact the motion of the body as they propagate [4]. This continuing influence of radiated waves is termed the memory effect because it makes the instantaneous forcing of the body dependent on its past motions.

Just as in the diffraction problem, a potential flow boundary value problem is set up to determine the forces in the radiation problem. In this case, a radiation potential, ϕ_R, is the velocity potential function required to describe the wave motion in the problem domain, with $\nabla^2\phi_R = 0$. The only difference in boundary conditions is that now the normal water velocity on the body surface must match the normal velocity of the surface caused by the body's velocity, $\mathbf{v}_{body}$.

$$\frac{\partial\phi_R}{\partial\hat{n}} = \mathbf{v}_{body} \cdot \hat{n} \quad \text{on} \quad S_{body} \tag{2.6}$$

As in the diffraction problem, the solution to this boundary value problem can be obtained using a panel method in the frequency domain, and the pressure calculated from the radiation potential solution can be integrated over the surface of the body to determine the net forces. It is important to note that, this time, the frequency in question is the frequency of the body oscillations and the generated waves – there are no incident waves. A seperate analysis is required for each DOF of the body. For each DOF and frequency, the resulting periodic forcing can be decomposed into two components, one in phase with the body acceleration (the added mass term) and one in phase with the body velocity (the damping term) [33]. For translational DOFs, this can be written as:

$$\mathbf{A}_{j,k} + i\mathbf{B}_{j,k} = -\rho \iint\limits_{S_{body}} \frac{\partial\phi_{Rj}}{\partial t}\ \hat{k} \cdot \hat{n}\ dS \tag{2.7}$$

where $\mathbf{A}_{j,k}$ and $\mathbf{B}_{j,k}$ are elements of six-by-six added mass and damping matrices, respectively, representing the force in DOF k from motion in DOF j. ϕ_{Rj} is the velocity potential from unit-amplitude oscillation of DOF j, and $\hat{k}$ is a unit vector in the direction of the DOF under consideration.

Superposition of Components

To summarize the coefficients produced by the three components of linear hydrodynamics:

1. The hydrostatic component represents the restoring force on the body by a frequency-invariant six-by-six stiffness matrix, $\mathbf{C}$.

2. The diffraction problem considers monochromatic waves (and a static body) in the frequency domain and yields a six-element, complex wave excitation force/moment vector normalized for wave amplitude, $\mathbf{X}(\omega, \beta)$, as a function of wave frequency and heading.

3. The radiation problem considers periodic oscillations of each DOF of a body (in still water) in the frequency domain and yields six-by-six added mass and damping ($\mathbf{A}(\omega)$ and $\mathbf{B}(\omega)$) matrices as a function of body oscillation frequency.

These coefficients can then be combined with the mass of the body, described by matrix $\mathbf{M}$, to form the hydrodynamic equations of motion. The generic form of the equation of motion for a floating body is:

$$(\mathbf{M} + \mathbf{A}(\omega))\ddot{\xi}(t) + \mathbf{B}(\omega)\dot{\xi}(t) + \mathbf{C}\xi(t) = Re\{A_{wave}\mathbf{X}(\omega, \beta)e^{i\omega t}\} \tag{2.8}$$

where $\xi(t)$ is a vector describing the system response in each DOF. The right hand side is the wave excitation force, where $\mathbf{A}_{wave}$ is the amplitude of waves with frequency ω and heading β.

These three components (one frequency-invariant, one dependent on wave frequency, and one dependent on body frequency) can be combined in one of two ways: by a frequency-domain or a time-domain approach. [1]

2.3.4 Frequency-Domain Linear Hydrodynamics

Frequency-domain analysis assumes that incident waves are monochromatic, that the equations of motion are linear, and that the body response has achieved steady state and therefore is of the same frequency as the incident waves. These assumptions limit the applicability of the approach because they preclude irregular wave conditions and non-harmonic body motions such as might occur from external forces (such as wind loads) or nonlinear system dynamics. However, because of its simplicity, a frequency-domain approach is useful for preliminary optimization work and for studying the hydrodynamics in a simplified form.

[1] For a dedicated comparison of the two approaches in the context of wave energy converter modelling, see [36].

Complex amplitudes are used to represent the frequency-domain response. For example, the complex amplitude vector for platform displacement, $\boldsymbol{\Xi}$, is defined such that

$$\xi = \boldsymbol{\Xi} e^{i\omega t} \tag{2.9}$$

$$\dot{\xi} = i\omega\ \boldsymbol{\Xi} e^{i\omega t} \tag{2.10}$$

$$\ddot{\xi} = -\omega^2 \boldsymbol{\Xi} e^{i\omega t}. \tag{2.11}$$

Making use of the assumptions that the waves and body motion are harmonic and of the same frequency, (2.8) can be rewritten as:

$$[-\omega^2(\mathbf{M} + \mathbf{A}(\omega)) + i\omega\mathbf{B}(\omega) + \mathbf{C}]\boldsymbol{\Xi}(\omega) = A_{wave}\mathbf{X}(\omega, \beta) \tag{2.12}$$

The complex response amplitudes, $\boldsymbol{\Xi}(\omega)$, can then be solved for if the coefficients are known. The vector of wave excitation coefficients, $\mathbf{X}$, is complex to account for phase. The frequency-dependent response, normalized by wave height, in terms of DOF amplitudes and phases, is commonly referred to as a response amplitude operator (RAO), where

$$RAO_k(\omega) = \frac{\boldsymbol{\Xi}_k(\omega)}{A_{wave}(\omega)} \tag{2.13}$$

with k specifying the degree of freedom. As such, translational RAOs are unitless and rotational RAOs have units of m^{-1}. An RAO is a convenient way to describe the steady-state frequency response of a floating structure's DOFs to wave excitation.

It is often useful to aggregate the frequency-dependent response across the frequency range into a single value that represents the overall magnitude of response when the system is subject to a spectrum of excitation. The metric that is often used in this case is the variance of the time series response; this is a convenient way to summarize the amount of energy in the aggregate response. Taking the square root to obtain the standard deviation or root-mean-square provides a measure of the "significant" motion amplitude, analogous to the notion of significant wave height used for irregular waves [37]. According to linear systems theory, for a linear system subject to a Gaussian-distributed random input (as is assumed for the wave elevation time series resulting from power spectral density $S(\omega)$), the output will also be Gaussian distributed, and its autocorrelation will equal its variance. Furthermore, the Wiener

Khinchin theorem states that the autocorrelation equals the Fourier transform of the absolute square of the output [38]. Thus the variance of the platform displacement in DOF k can be calculated as

$$\sigma_{\xi_k}^2 = \int_0^\infty |\Xi_k(\omega)|^2 d\omega \tag{2.14}$$

where the complex amplitude, $\Xi_k(\omega)$, is equivalent to the Fourier transform of the response time series, $\xi_k(t)$. Alternatively, the same quantity can be calculated based on the response amplitude operator, RAO_k, and the wave spectrum, $S(\omega)$, in which case (2.14) becomes:

$$\sigma_{\xi_k}^2 = \int_0^\infty |RAO_k(\omega)|^2 S(\omega) d\omega \tag{2.15}$$

In this way, a harmonic response calculated in the frequency domain can be used to estimate the characteristic amplitude and energy of the response in the time domain.

2.3.5 Time-Domain Linear Hydrodynamics

The time-domain linear hydrodynamics approach has a number of important advantages over the frequency domain approach, including:

- external loads and kinematic constraints can be applied without restriction, permitting coupling with other systems
- transient analysis is possible – the platform does not have to oscillate at the incident wave frequency
- irregular wave environments can be modelled by superimposing any number of regular (monochromatic) waves

While the coefficients calculated from the diffraction and radiation problems (which are already in frequency-domain form) combine effortlessly into a frequency-domain model, putting them into a time-domain model is less straightforward.

Diffraction in the Time Domain

The frequency-dependent wave excitation coefficients, $\mathbf{X}(\omega)$, calculated from the diffraction problem need to be applied to a wave elevation time series, $\zeta(t)$, to produce an excitation force time series, $\mathbf{f}_e(t)$.

In many cases, the wave elevation time series for the simulation is generated from a frequency-domain spectrum, $S(\omega)$, so a frequency-domain representation of the incident waves is already available. (Alternatively, if a wave elevation time series recorded from a wave buoy or time-domain wind-wave model is used, a Fourier transform could be applied to yield the desired frequency-domain wave representation.) This frequency-domain representation can be multiplied by the frequency-dependent excitation coefficients to obtain the excitation force in the frequency domain. This can then be inverse Fourier transformed to obtain the wave excitation force time series:

$$\mathbf{f}_e(t) = \frac{1}{2\pi}\int_0^\infty W(\omega)\sqrt{2\pi S_\zeta(\omega)}\mathbf{X}(\omega,\beta)e^{i\omega t}d\omega \tag{2.16}$$

where $W(\omega)$ is the Fourier transform of a white Gaussian noise time series process with zero mean and unit variance [4]. It is complex, and used to select the amplitudes and phases of the discrete realization of the wave spectrum.

Alternatively, the excitation force time series could be calculated in the time domain, by taking the convolution of wave height and the wave excitation force kernel (or wave excitation impulse response function), $\mathbf{k}_e(t)$:

$$\mathbf{f}_e(t) = \int_{-\infty}^\infty \zeta(\tau)\mathbf{k}_e(t-\tau)d\tau \tag{2.17}$$

The kernel $\mathbf{k}_e(t)$ represents the forcing response to a Dirac delta function in the wave elevation time series and is calculated by the inverse Fourier transform of the wave excitation coefficient [4]:

$$\mathbf{k}_e(t) = \frac{1}{2\pi}\int_{-\infty}^\infty \mathbf{X}(\omega,\beta)e^{i\omega t}d\omega \tag{2.18}$$

Radiation in the Time Domain

Application of the frequency-dependent added mass and damping terms generated from the radiation problem in the time domain follows a similar approach to that of equation (2.17), in order to capture the frequency dependence of the coefficients and the resulting memory effect as discussed in Section 2.3.1. Because the body motion is known in the time domain (and accurate conversion of the instantaneous body motions to the frequency domain would require future knowledge), the calculation is best done in the time domain, using the approach of a convolution integral.

As with the method of equation (2.17), the frequency-domain added mass and damping coefficients need to be converted into time-domain form, this time in the form of a wave-radiation retardation kernel (or radiation impulse response function), k_r, representing the transient force response to a unit impulse in body velocity [33]:

$$k_r(t) = \frac{2}{\pi}\int_0^\infty Bcos(\omega t)d\omega \tag{2.19}$$

Here, $k_r(t)$ and $B(\omega)$ are six-by-six matrices, with matrix position i, j corresponding to the force or moment in DOF i resulting from movement in DOF j. With the assumption of linearity, a convolution integral with k_r and the body velocity vector, $\dot{\xi}$, can be used to obtain the instantaneous forcing caused by wave-radiation damping [4]:

$$f_r(t) = \int_0^t \mathbf{k}_r(t-\tau)\dot{\xi}(\tau)d\tau \tag{2.20}$$

As the Kramers-Kronig relation shows, this convolution integral also includes the memory effect, or the frequency dependence, of added mass forces. The Kramers-Kronig relation relates added mass and damping through the following identity, which is found in many linear hydrodynamics formulations [4, 33, 36]:

$$\mathbf{A}_{i,j}(\infty) = \mathbf{A}_{i,j}(\omega) + \frac{1}{\omega}\int_0^\infty \mathbf{k}_{r\,i,j}(t)sin(\omega t)dt \tag{2.21}$$

where $\mathbf{A}_{i,j}(\omega)$ is frequency-dependent added mass and $\mathbf{A}_{i,j}(\infty) = \lim_{\omega\to\infty}\mathbf{A}_{i,j}(\omega)$ is infinite-frequency added mass. This shows that even though $\mathbf{k}_r$ is calculated based solely on damping coefficient, $\mathbf{B}(\omega)$, equation (2.20) includes the parts of added mass that provide a memory effect; damping and added mass coefficients are interrelated. The remaining portion of the added mass in equation (2.21), the infinite-frequency added mass, is frequency-independent and can be added as a constant to the body's mass matrix [36].

Time-Domain Superposition of Components

Combining the three components of linear hydrodynamics in the time domain gives the following matrix equation of motion:

$$(\mathbf{M} + \mathbf{A}_\infty)\ddot{\xi}(t) + \int_0^\infty \mathbf{k}_r(t-\tau)\dot{\xi}(\tau)d\tau + \mathbf{C}\xi(t) = \mathbf{f}_e(t) \tag{2.22}$$

where $\mathbf{f}_e$ can be calculated in the method of equation (2.16) or (2.17).

Limitations of Time-Domain Linear Hydrodynamics

While linear hydrodynamics is a powerful and widely-used method for modelling floating structure dynamics, linear hydrodynamics theory as described above neglects higher-order inviscid effects as well as viscous effects that can be significant.

Second-order inviscid effects arise from changes in wave excitation, added mass, and damping forces when the platform is displaced from its equilibrium position in the water. One such effect, not captured by linear hydrodynamics, is the mean-drift force in which wave loads result in a net force on the body in the direction of wave propagation. Another second-order effect is sum-frequency excitation, which is excitation occurring at frequencies of the sum of individual wave component frequencies. These can be a particular problem for TLPs, which have very high natural frequencies in heave with very little damping. These natural frequencies are well above the energetic wave frequencies, but not above the summed frequencies [39].

The assumptions of non-steep waves further limits linear hydrodynamics to mild wave conditions. Steep and breaking waves, which become more likely in shallow waters or extreme weather conditions and can exert large transient forces on a platform, cannot be modelled by linear hydrodynamics.

Neglecting viscous drag is another weakness with linear hydrodynamics in certain situations. Some platform designs make use of heave plates specifically meant to create large viscous drag forces to damp heave motion. For these sorts of designs, the inclusion of a viscous drag model is essential. A common solution is to use the approach of Morison's equation to add a viscous drag term to the overall dynamics equations. Although Morison's equation is meant for slender cylinders, it can work as a useful approximation of viscous drag to an otherwise-inviscid model if appropriate drag coefficients are chosen, as is done in the models discussed in [4] and [6].

2.3.6 Higher-Fidelity Hydrodynamics Treatments

The Eulerian nature of typical CFD methods – where the fluid domain is discretized by a mesh – is challenged by the free surface involved in the hydrodynamics of a floating structure. Representing and modelling the free surface within such a discretization requires a special approach. The most common approach is the volume of fluid (VOF) approach, in which the free surface boundary is accounted for by a single variable in

each cell describing the fractional volume of fluid within the cell [40]. A practical disadvantage of the VOF approach is that it adds additional computation to the already-daunting computation requirements of CFD.

Smoothed-particle hydrodynamics, a Lagrangian approach in which the fluid is modelled by particles that move with the fluid velocity and interact according to the Navier-Stokes equations, is well-suited to modelling the free surface and even phenomena such as breaking waves and flow separation. Unfortunately, at present, the large number of particles required for accuracy makes smoothed particle hydrodynamics quite computationally-intensive [41].

2.3.7 Current Trends

Linear hydrodynamics and strip theory, the two common modelling approaches for platform hydrodynamics, to a large degree account for different sets of hydrodynamic loads. If the added mass term in strip theory is removed (because it is already included in linear hydrodynamic theory), it can be argued that the two theories are complimentary; they model independent phenomena and do not result in any sort of double counting when applied together. (While it is true that in some cases a strip theory approach is more suitable and in others a linear hydrodynamics approach is more suitable, in either case, the reason the other approach is less suitable is that the phenomena it models are of diminished importance – therefore, its inclusion has minimal impact on the simulation.) The combination of both approaches is therefore advisable. While the assumptions of linear hydrodynamics make its use for large platform motions questionable, the wave radiation phenomena neglected by strip theory are crucial for many geometries, so there is little merit in using just one of these approaches. The limiting assumptions of linear hydrodynamics are hard to avoid without moving to slower, higher-fidelity modelling techniques.

Higher-fidelity approaches such as volume of fluid methods and smoothed-particle hydrodynamics can do away with the limitations of linear hydrodynamics and strip theory and provide more detailed results. However, the computational cost associated with such improvements is well beyond the norms for floating wind turbine design codes, so these modelling approaches are not suited for use in design studies at present.

2.4 Mooring Line Dynamics Modelling

Mooring lines serve a station-keeping function – restraining the platform from drifting in the water under current, wave, and wind loads. In floating turbine configurations with yaw drives, the mooring lines fix the platform's yaw DOF so that the rotor yaw can be positively controlled. The mass and hydrodynamic drag forces on the mooring lines also contribute to the damping of the platform's motions. For TLPs, taut mooring lines also provide the primary means of stability for the platform.

Mooring line models range from simple linear or nonlinear six-DOF force-displacement relationships to quasi-static numerical approaches and fully-dynamic FEM approaches with varying levels of complexity. Different approaches are suited to different types of modelling. An illustration of a mooring line, demonstrating the anchor and fairlead locations, is given in Figure 2.5.

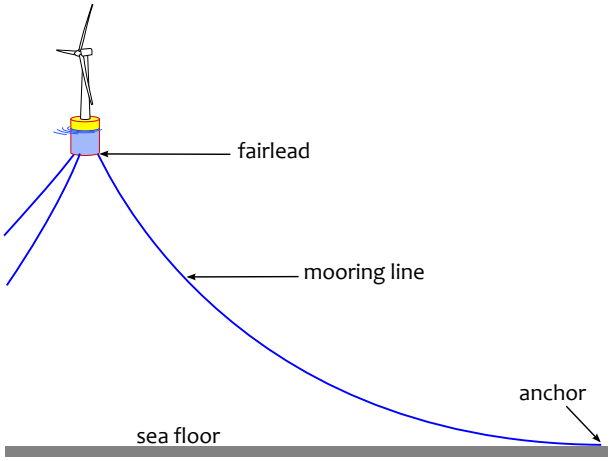

Figure 2.5: Mooring line anatomy

2.4.1 Force-Displacement Models

Force-displacement models are the simplest mooring line models. They make use of functional relationships between platform displacements/rotations and the resulting forces/moments from the mooring lines. As such, they ignore any dynamic effects on the lines. They can model the mooring system as a whole by operating in the six DOFs of the platform, or they can model each mooring line individually by operating on the three translational DOFs of each fairlead location independently. Linear versions without coupling between DOFs consist of just an effective stiffness in each DOF. For

a frequency-domain model of the overall system, a linear force-displacement mooring model is necessary in order to maintain the linearity required for the system response to be time-harmonic.

2.4.2 Quasi-Static Models

Quasi-static models capture some of the nonlinear behaviour of mooring lines, and can include effects like sea bed friction and axial stiffness. However, they neglect the motion of the mooring lines and the surrounding water (which may be moving from waves or currents). This means that the inertia of the lines and drag forces on the lines are excluded from the model; the mooring system provides a stiffness but not a damping effect in simulations using these models.

With only tension and its own weight acting on it, a mooring line will lie in a vertical plane with opposite corners defined by the fairlead and anchor locations, and it will sag in the shape of a catenary, which can be described analytically. With the fairlead and anchor points defined (as well as cable weight, stiffness, and seabed friction) the tensions at the fairlead end of the cable and subsequently the tensions and positions throughout the cable can be obtained by numerically solving a system of two nonlinear equations, making for a relatively fast computation [4].

2.4.3 Dynamic Models

Dynamic models use a discretization approach to describe the motion of the mooring lines, and the forces on them, in terms of the dynamics of a concatenation of simple segments. This is usually done with a finite element method, though a finite difference method can also be used. In addition to the added details of hydrodynamic drag, added mass, and cable inertia, additional cable properties such as bending and torsional stiffness can be included.

In a dynamic model, the cable no longer necessarily assumes a catenary shape nor lies in a vertical plane. The cable also has velocities along its length which must be calculated. These are additional variables that simulators using fully dynamic cable models need to support. In high-fidelity models, seabed friction will play a roll in both axial stretching and lateral rolling/dragging across the seabed.

While a dynamic model is necessary to account for the added inertial and damping effects of the mooring lines on the floating platform dynamics, the computational

expense of dynamic models has limited their use in coupled floating wind turbine design codes.

2.4.4 Current Trends

Though modelling of the mooring lines is usually given secondary importance compared to the wind turbine and floating platform models, a range of well-established mooring line modelling options exist. Simple force-displacement models are uncommon because quasi-static models are vastly superior and are still quite fast thanks to their analytical formulation. Because they capture all of the effects that make up the primary function of the mooring system – station keeping – quasi-static models are extremely popular. However, the validity of neglecting the dynamic effects of the mooring lines is an issue that has not been thoroughly explored in the past [29]. The improved accuracy of dynamic mooring line models is countered by their computational expense. Whether their improved accuracy is necessary in some cases is the question Chapter 3 of this thesis seeks to answer.

2.5 Third-Party Models Used in This Thesis

Three preexisting modelling tools were used extensively in the work of this thesis: the coupled time-domain floating wind turbine simulator FAST, the time-domain cable dynamics simulator ProteusDS, and the frequency-domain linear hydrodynamics preprocessor WAMIT.

2.5.1 FAST

Developed by the US National Renewable Energy Laboratory (NREL), FAST is one of the three most prevalent options for coupled floating wind turbine simulation. The other two are HAWC2, by Risø DTU, and GH Bladed, by Garrad Hassan. FAST stands our from the other two options by being open-source and having a more advanced platform hydrodynamics modelling capability.

FAST is a nonlinear time-domain simulator that combines limited structural degrees of freedom with BEM aerodynamics, linear hydrodynamics, and quasi-static mooring dynamics in order to model most conventional horizontal-axis wind turbine configurations. A linearized representation of the aero-hydro-servo-elastic system can also be generated by FAST for use in controls development.

FAST uses a linear mode shape approach to model the structural deflections of the tower and rotor blades. It uses two flapwise DOFs and one edgewise DOF in each blade, and two fore-aft DOFs and two lateral DOFs in the tower. Additional DOFs include the generator rotation, nacelle rotation, a spring-damper model of the low speed drive shaft torsion, and a full six degrees of freedom for the base of the turbine to account for movement of a floating platform. The entire formulation of the tower and turbine assumes small tilt angles of the support platform [4].

To achieve it's full aero-hydro-servo-elastic functionality (combining the effects of aerodynamics, hydrodynamics, controls, and structural dynamics), FAST interfaces with an aerodynamics subroutine, AeroDyn, and a hydrodynamics subroutine, HydroDyn, for offshore installations [4].

AeroDyn uses a BEM model or a generalized dynamic wake model for the rotor aerodynamics. It includes tip and hub loss corrections and a Beddoes-Leishman dynamic stall model. The aerodynamics are fully coupled to the motion of the blades [4]. AeroDyn uses a longer time step than FAST as the aerodynamics have longer time constants than the structural dynamics.

HydroDyn treats the platform as a rigid body and assumes its center of mass and center of buoyancy are along the centerline of the turbine tower. The hydrodynamics on the platform are composed of the three components of linear hydrodynamic loadings, as described in Section 2.3.3, as well as the viscous drag term from Morison's equation. The linear hydrodynamic coefficients are calculated using a frequency-domain preprocessor. The viscous drag term augments the linear hydrodynamics model with an approximated viscous drag force, calculated using strip theory and an "effective diameter" of the platform. This viscous drag treatment does not take into account the motion of the platform or the instantaneous free surface height in order to be more consistent with the linear formulation of the rest of the hydrodynamics problem [4]. The hydrostatic restoring force is treated with a six-by-six stiffness matrix. Wave excitation is treated in the method of equation (2.16) to calculate the total wave force from irregular waves at each instant in time. HydroDyn models irregular waves using a JONSWAP or custom-defined wave spectrum. Nonlinear waves or waves coming from multiple directions are not supported. Wave-radiation damping is treated using convolution of a wave radiation retardation kernel as in equation (2.20).

FAST features a quasi-static mooring line model; the positions of the cables are solved for at each time step assuming that each cable is in static equilibrium in still water, thereby neglecting cable inertia and hydrodynamic damping, as discussed in

2.4.2. The model includes cable weight, axial stretching, and seabed friction, but does not include bending stiffness. The quasi-static approach is rationalized based on the small proportion of system mass made up by the cables, and it is pointed out that ignoring cable damping is a conservative assumption [4]. Testing this rationalization is the focus of Chapter 3 of this work.

A number of replacement hydrodynamics and mooring line models have been previously used with FAST, including Charm3D, which adds a FEM-based mooring line model [29], and TimeFloat, which adds viscous damping modelling of heave plates and other platform elements using an adapted version of Morison's equation and special line members with drag coefficients calibrated from experimental data [6].

FAST is the model used for the coupled floating wind turbine simulations in Chapter 3. It is also the model used to generate the linearized wind turbine coefficients for the frequency-domain modelling in Chapters 4 and 5.

2.5.2 ProteusDS

ProteusDS is the dynamic mooring line model that is used in Chapter 3 to evaluate the limitations of FAST's built-in quasi-static mooring line model. It is a time domain hydrodynamics simulator, developed by Dynamic Systems Analysis Ltd. and researchers at the University of Victoria, specializing in underwater cable and net dynamics.

ProteusDS's rigid body hydrodynamics models uses a combination of Morison's equation and a Froude-Krylov force approach for wave excitation and hydrostatics, taking into account the variable wetted area of the platform. As such, the hydrodynamics include both viscous and pressure sources of loading from the undisturbed waves and the body motion, but do not include diffraction or radiation effects of the body on the water. Alternatively, kinematic RAOs can be used for the body response, however such kinematic RAOs preclude coupling with external forces or transient analysis.

ProteusDS features a dynamic nonlinear FEM-based cable model that makes use of 3rd-order cable elements with bending and torsional stiffnesses. Its high-fidelity cable model is why it was chosen for use in Chapter 3. The ProteusDS simulator is available in DLL form, which enables coupling to wind turbine modelling codes to provide coupled floating wind turbine simulations making use of its high-fidelity mooring line model.

Two well-established alternatives to ProteusDS are Charm3D, developed by Texas A&M University, and Orcaflex, developed by Orcina. Both of these simulators feature a similar FEM mooring line model to ProteusDS, and a more advanced floating structure hydrodynamics capability that includes first- and second-order hydrodynamics calculated using WAMIT [29, 42].

2.5.3 WAMIT

WAMIT (WaveAnalysisMIT) is the linear hydrodynamics solver that is used to provide the input coefficients to FAST's hydrodynamics model, HydroDyn, in Chapter 3 and the frequency-domain hydrodynamics models used in Chapters 4 and 5. Developed by researchers at MIT and first released in 1987, it is the industry standard for linear hydrodynamics preprocessing. WAMIT uses a numerical panel method to find the harmonic solutions of the diffraction and radiation boundary value problems for rigid bodies in water. Being based on linear hydrodynamics, WAMIT models the fluid domain using inviscid potential flow theory and uses Airy wave kinematics. Being a panel-method solver, it discretizes the body surfaces into rectangular panels and then solves for an unknown potential source strength at each panel such that the boundary conditions (described in Section 2.3.3) are satisfied [35].

2.6 Modelling Summary

The nature of floating wind turbine modelling is such that a number of different models for the different parts of the structure and the physical phenomena acting on them need to be coupled together to provide an overall simulation of the system. To evaluate the accuracy of a simulation, the models need to be studied in combination rather than independently. There are a range of options for each modelling area, with different levels of fidelity and corresponding computational costs. In each modelling area (turbine, platform, and mooring lines), computationally-efficient lower-fidelity models are the current norm, but higher-fidelity alternatives exist that are below the computational burden of the highest-fidelity approaches but still offer significant improvements of the currently-standard methods. The adequacy of all three modelling areas is a subject of ongoing research, and the adequacy of the mooring models is particularly under-explored. Several well-established coupled simulation capabilities exist, as well as higher-fidelity models that can be coupled into them. However, these

simulators currently focus on using lower-fidelity methods and each have modelling limitations that are driving research efforts toward their improvement.

Chapter 3

Evaluating the Adequacy of Quasi-Static Mooring Line Models

3.1 Introduction

Floating wind turbine simulators currently tend to use quasi-static mooring line models in the form of either force-displacement relationships or analytical solutions for catenary cables in static equilibrium [29]. These approaches have the advantage of computational efficiency, which is desirable since the structural, hydrodynamic, and aerodynamic models commonly used also tend to run quickly. For cases where waves are small, and support platform and mooring line velocities are minimal, quasi-static models can provide a good approximation to the true system dynamics. For cases with higher platform and mooring line velocities, dynamic effects neglected by quasi-static mooring models may be significant. As different support structure stability classes have different amounts of motion, the adequacy of quasi-static mooring models depends on the stability class being evaluated. Often, the primary effect of the mooring line dynamics on the overall system is an increase in the damping on the platform, which benefits platform stability. For that reason, it has been argued that using a quasi-static model under-predicts the stabilizing effect of the mooring lines and is therefore a conservative modelling approach [4].

However, this conservative approach has drawbacks. In the offshore wind industry, where the profit margins are slimmer than in other offshore industries such as oil and gas, the overdesign resulting from under-predicting the damping on the platform may compromise competitiveness. Offshore wind turbines, unlike oil platforms,

are unmanned so safety factors can be lower and conservative design is not as critical. Moreover, the mooring line inertia and hydrodynamics that are ignored in a quasi-static model can produce additional system motions that can lead to increased structural loading on the turbine – snap loading of the mooring lines being a prime example.

Absent from the literature is a broad investigation into the importance of dynamic effects of the mooring lines for a given system. Cordle identified the need for a dedicated study into the importance of dynamic mooring line effects for floating wind turbines, and how this sensitivity might vary for different water depths and system designs [29]. A number of recent studies make steps in that direction. Waris and Ishihara [43] compared coupled simulation results of a three-column semisubmersible design using a linear force-displacement mooring model and a dynamic finite element mooring model, as well as experimental results. Their study identified important limitations to linear force-displacement models but did not consider nonlinear quasi-static models, which are a practical middle ground between the options they considered. A study by Kallesoe and Hansen [12] used the simulation code HAWC2 with the addition of an FEM-based dynamic mooring model. The original, quasi-static mooring model in HAWC2 uses lookup tables describing the nonlinear force-displacement relationships of the mooring system. The study analyzed the dynamics of the OC3 Hywind system under a number of normal operation conditions, and concluded that the dynamic mooring model showed similar extreme loads but reduced fatigue loads compared to the original quasi-static model. Matha et al. also compared a dynamic model with a quasi-static model on the OC3 Hywind system, this time using a multibody mooring line model which they developed [44]. They observed that the dynamic model resulted in reduced motion amplitudes and increased natural periods of the system. Dynamic mooring models have also been developed for the FAST simulator. Azcona et al. developed one such model, featuring a finite-element mooring line model with axial elasticity, nonlinear hydrodynamics, and bottom contact dynamics [45]. Another option is the finite-element OrcaFlex model, with a new module for coupling with FAST; Masciola et al. used this to compare OrcaFlex with FAST's built-in quasi-static mooring model [42]. Their study focussed more on the numerical stability and functionality of the FAST-OrcaFlex coupling rather than the significance of the dynamic effects introduced by the OrcaFlex model. The study presents a sample of results using the OC3 Hywind, but mentions that the ITI Energy Barge and MIT/NREL TLP were also studied. Numerical instabilities were encountered when

simulating the taut lines of the TLP design. Though these studies provide valuable contributions in model development and model evaluation on specific system designs, a more extensive study comparing dynamic and quasi-static mooring models across multiple classes of floating support structure is still missing.

The work presented here provides a more thorough look at the importance of mooring line dynamics by using a more advanced floating wind turbine simulator, using more test scenarios, and considering several different support structure configurations. The coupled floating wind turbine simulation code used is FAST, which features a quasi-static mooring model that solves a set of analytical catenary cable equations to determine the instantaneous static-equilibrium positions and tensions of the mooring lines. This quasi-static model is compared with a dynamic FEM-based mooring model, referred to as ProteusDS, that has been coupled with the FAST simulator. FAST remains in charge of the floating platform hydrodynamics modelling in both cases. Three standard turbine design concepts – the OC3-Hywind, the ITI Energy Barge, and the MIT/NREL TLP – are analyzed in the comparison effort across a range of test conditions. These test conditions vary from carefully-controlled still-water tests that can isolate the damping contributions of the mooring lines, to tests in turbulent winds and irregular waves that can reveal how much the mooring line dynamic effects affect the overall system response in realistic operating conditions. The three design concepts represent the three different stability classes. Simulating identical scenarios with both the quasi-static and the dynamic mooring models allows detailed evaluation of the impact of the mooring line dynamics on the overall simulation.

The interface created to couple ProteusDS to the FAST simulator, the three floating turbine designs tested, and the test cases they were simulated in are described in section 3.2. Section 3.3 presents the results of the comparison effort, showing the static equivalence of the two mooring models, and how they differ in damping the platform motions and affecting the response of the systems under various conditions. Section 3.5 reiterates the main findings from the comparison effort and identifies remaining research work to be done in this area.

3.2 Methodology

3.2.1 Coupled Simulator

The time-domain floating wind turbine simulator FAST, which was described in Section 2.5.1, features a computationally-efficient quasi-static mooring line subroutine, Catenary, that models taut or slack catenary mooring lines. It accounts for weight and buoyancy, axial stiffness, and friction from variable contact on the seabed, but does not account for bending or torsional stiffnesses. Being quasi-static, it solves for cable positions and tensions under static equilibrium given the instantaneous fairlead (cable-platform connection) location. Cable inertia and hydrodynamic forces are ignored [4]. This makes each mooring line lie in a catenary shape along a vertical plane whose corners are defined by the fairlead and bottom-contact locations. With this simplification, the horizontal and vertical tensions at the fairlead can then be solved for by the numerical solution of two analytical equations of those two unknowns. If there is no seabed contact (the angle of the line at the anchor is greater than zero), then the following equations are used.

$$x_F(H_F, V_F) = \frac{H_F}{\omega}\left\{\ln\left[\frac{V_F}{H_F} + \sqrt{1+\left(\frac{V_F}{H_F}\right)^2}\right] - \ln\left[\frac{V_F-\omega L}{H_F} + \sqrt{1+\left(\frac{V_F-\omega L}{H_F}\right)^2}\right]\right\} + \frac{H_F L}{EA} \quad (3.1)$$

$$z_F(H_F, V_F) = \frac{H_F}{\omega}\left[\sqrt{1+\left(\frac{V_F}{H_F}\right)^2} - \sqrt{1+\left(\frac{V_F-\omega L}{H_F}\right)^2}\right] + \frac{1}{EA}\left(V_F L - \frac{\omega L^2}{2}\right) \quad (3.2)$$

where x_F and z_F are the horizontal and vertical fairlead coordinates, H_F and V_F are the horizontal and vertical components of the line tension at the fairlead, ω is the line weight per unit length, L is the unstretched line length, and EA is the axial stiffness of the mooring line. A modified set of equations is used when part of the cable length is in contact with the sea floor [4].

These equations provide the forces from the mooring lines. FAST uses additional analytic equations to specify the positions and tensions along the length of each line.

3.2.2 Dynamic Mooring Model

The dynamic mooring model ProteusDS was introduced in Section 2.5.2. Given the lack of experimental work on floating wind turbines, this model, integrated into a floating wind turbine simulation, can provide an accurate reference to compare

results from a quasi-static mooring line model against. The cable model of ProteusDS has been validated with experimental results [46]. It can simulate floating systems composed of both rigid structures and flexible cables and nets. Its customizable joint capabilities allow for the simulation of mooring lines with segments having different physical properties, as well as mooring systems featuring floats, clump weights, and delta connections. ProteusDS features a dynamic nonlinear cable model, based on the work of Buckham et al. [47]. In addition to the effects included in a quasi-static model, this model includes the effects of the distributed hydrodynamic loading on the cable, the internal shear and axial forces within the cable, and the friction and restitutional forces that develop from contact with the seabed. The continuous form of the equation of motion at any point on a mooring cable is [47]:

$$-(EI\mathbf{r}'')'' + [((EA\varepsilon + C_{ID}\dot{\varepsilon}) - EI\kappa^2)\mathbf{r}'] + [GJ\tau(\mathbf{r}' \times \mathbf{r}'')]' + \mathbf{h} + \mathbf{w} = \mathbf{M}_I\ddot{\mathbf{r}} \quad (3.3)$$

where $\mathbf{r}(s,t)$ is the absolute position of the cable centerline with components defined relative to an inertial reference frame, $\mathbf{h}$ is the hydrodynamic load per unit length, $\mathbf{w}$ is the apparent weight of the cable per unit length, $\mathbf{M}_I$ is a 3 by 3 mass matrix (specific to the inertial frame's axes) that includes direction-dependent added mass terms, and C_{ID} is a damping coefficient representing axial structural damping in the cable. Variables ε, κ, and τ are the axial strain, curvature, and twist, respectively, and EA, EI, and GJ are the respective stiffnesses of the cable for each type of deformation.

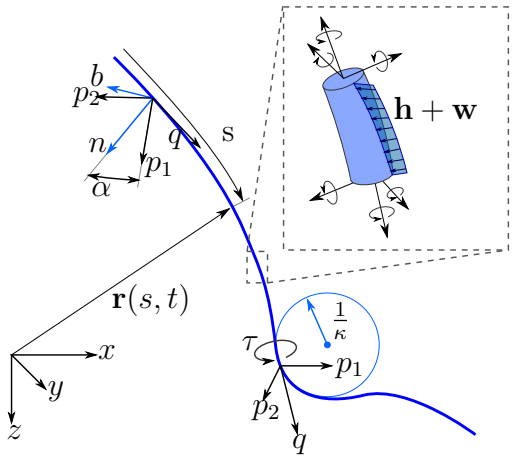

Figure 3.1: Coordinate systems of the ProteusDS mooring line model

ProteusDS uses a finite-element discretization of equation (3.3) with the cable

mass lumped at the node points. The state of a simulated cable is defined by a cubic spline approximation to the cable centerline, $\mathbf{r}(s,t)$, plus a linear approximation to the twist, τ, of the cable cross section about the centerline. The state variables used in these two polynomial approximations are absolute positions, curvatures, and a rotation angle, α, to account for the cable twist at each node point. By expressing the cable dynamics in terms of absolute coordinates, the model is a form of the Absolute Nodal Coordinate Formulation (ANCF) presented by Berzeri and Shabana [48]. The twisted spline approximation ensures second order continuity between elements but, to ensure smoothness in the assembly, boundary conditions must be applied between elements. These additional constraint equations allow the curvatures to be recovered at any time from the assembled set of node positions. Thus, for an assembled cable, the curvatures can be eliminated from the discretized equations.

The cable twist model is composed of a torsion calculated from the approximated centerline and the rotation angle, α, which describes the rotation of the cross section about the centerline. This angle is measured relative to the Frenet frame, defined by normal vector n and binormal vector b, which can be recovered from the cable centerline profile, $\mathbf{r}(s,t)$. Thus, given a set of node positions and a series of cross section rotations, the axial, bending, and torsional strains experience in the cable can be calculated. Equation (3.3) is used to explicitly solve for node accelerations that are integrated to determine the complete cable configuration at the next time step [47].

The hydrodynamic forces are calculated using Morison's equation. The added mass contribution – a function of the cable volume, orientation, water density, and an added mass coefficient, C_a – is included in $\mathbf{M}_I$. The drag term, $\mathbf{h}$ in equation (3.3), is calculated as:

$$\mathbf{h} = -\frac{1}{2}\rho d_C \mathbf{R}_I \begin{bmatrix} f_p C_D \frac{-v_{p1}}{\sqrt{v_{p1}^2+v_{p2}^2}} |\mathbf{v}|^2 \\ f_p C_D \frac{-v_{p2}}{\sqrt{v_{p1}^2+v_{p2}^2}} |\mathbf{v}|^2 \\ -f_q C_D \operatorname{sgn}(v_q) |\mathbf{v}|^2 \end{bmatrix} \tag{3.4}$$

where $\mathbf{R}_I$ is a rotation matrix from a reference frame attached and oriented with the cable cross section relative to the inertial reference frame, d_c is cable diameter, and C_D is a normal drag coefficient, set to 2.0 in the work presented here. The loading coefficients, f_p and f_q, account for the distribution of hydrodynamic drag between the normal and tangential directions. They are non-linear functions of the angle

between the cable centerline and the relative velocity vector, as discussed by Driscoll and Nahon [49]. The cable's relative velocity is given by $\mathbf{v} = \dot{\mathbf{r}} - \mathbf{j}$, where $\mathbf{j}$ is the water's absolute velocity at the point considered, calculated in this case using Airy wave theory. When expressed in the local frame, the components of $\mathbf{v}$ are v_{p1}, v_{p2} and v_q where $p1$, $p2$ and q are the normal, binormal and tangent directions.

An added mass coefficient of 0.5 is applied for cable accelerations in the normal direction in the present work. The seabed contact forces are modelled using a constant friction coefficient for sliding contact forces (0.2) and linear stiffness (1000 kN/m) and damping (100 N-m/s) coefficients for the normal contact forces.

The ProteusDS cable model is available in dynamic link library (DLL) form, which enables coupling to aerodynamic or aero-servo-elastic wind turbine codes for coupled floating wind turbine simulation making use of ProteusDS's dynamic cable model.

3.2.3 FAST-ProteusDS Coupling

The DLL form of the ProteusDS simulator allows coupling with FAST. The coupling between the two simulators was arranged such that FAST transmits fairlead locations and velocities to ProteusDS, and ProteusDS in turn provides the mooring line reaction forces at the fairleads, as well as tensions and positions at the line nodes, to FAST. This communication is performed at every FAST time step. With its FEM implementation, the ProteusDS model requires a much smaller time step – approximately one-tenth the FAST time step. Though ProteusDS supports a variable time step, a fixed time step was used to prevent recurring hang-ups in the simulation. Averaging of mooring line tensions and positions to match between the two different time steps is handled internally by the ProteusDS DLL. Wave kinematics are synchronized between the models; the discretized seastate precalculated by FAST at initialization is automatically written to a ProteusDS input file. ProteusDS then uses these wave parameters to calculate the water velocities along the mooring lines at each time step.

A number of changes were made to the FAST source code to couple ProteusDS into the platform dynamics subroutine (HydroDyn). These include the addition of a new subroutine that interacts with the ProteusDS DLL in place of the default quasi-static mooring subroutine, and modifications to the FAST input and output variables to carry additional mooring system information and streamline switching between the two mooring models. In addition, an intermediate DLL was created to resolve data type incompatibilities between FAST and ProteusDS and to provide

conversion between the platform DOFs and the positions of the fairleads. This DLL sends precalculated seastate information from FAST to ProteusDS, and also loads and stores information about the mooring system needed for initial condition generation and mooring fairlead position and force calculations. With this setup, the ProteusDS mooring line simulator can be used in FAST simulations with only minor additions to the input files. As such, it provides a coupled simulation capability that enables convenient comparison of quasi-static and dynamic mooring line behaviour.

3.2.4 Turbine System Descriptions

The three floating wind turbine concepts selected for this comparison using the ProteusDS dynamic mooring line model are the OC3-Hywind, the ITI Energy Barge, and the MIT/NREL TLP.

The OC3-Hywind is a ballast-stabilized spar-buoy design with three slack catenary mooring lines. It is based on the Statoil Hywind prototype that is the world's first MW-scale floating wind turbine, but modified for the purposes of the Offshore Code Comparison Collaboration (OC3) project, an international effort to compare and validate leading floating wind turbine simulation codes. The OC3-Hywind design was used to compare simulators in terms of aerodynamic, structural, and especially hydrodynamic models in the OC3 Phase IV study [23].

The ITI energy barge is a buoyancy-stabilized rectangular barge concept designed at the University of Glasgow and University of Strathclyde for ITI Energy. It has some ballast to provide a 4 m draft, in order to minimize wave slamming. Eight slack catenary mooring lines spreading off at 45 degree angles from each other are connected in pairs at the four corners of the barge [50].

The MIT/NREL TLP is a mooring-stabilized design based on the parametric optimization work of Tracy [51] at MIT with modifications made by Matha [52] at NREL to correct for some underpredicted pitch and roll inertias [50]. It features a partially-submerged cylindrical platform and eight taut vertical mooring lines connected in pairs to four spokes extending from the bottom of the platform. The cylinder is ballasted with concrete to provide enough stability in mild conditions for float-out before the mooring lines are installed.

These three design concepts, illustrated in Figure 1.2, have been implemented by NREL in their FAST code. Each design features the NREL 5 MW offshore reference wind turbine, a hypothetical design with standardized specifications to enable accu-

rate comparison of different wind turbine modelling tools. FAST input files for the three designs were available from NWTC. Selected properties of the designs are listed in Table 3.1. These properties come from reference [50] except for the mooring line EI and GJ values – these bending and torsional stiffnesses were not included in the specifications of any of the designs but are necessary for the ProteusDS model. They were estimated by scaling typical ratios of axial and bending/torsional stiffness found in published wire rope test results [53]. This approach is an approximation, but tests with the ProteusDS model showed negligible sensitivity to even order-of-magnitude changes in the cable bending and torsional stiffness values.

Table 3.1: Selected turbine system specifications

	OC3	ITI	TLP
Draft (m)	120	4	47.89
Total mass (kg)	7 466 000	5 452 000	8 600 000
Displacement (m^3)	8 029	6 000	12 180
Platform diameter (m)	6.5 m	40x40	18
Water depth (m)	320	150	200
Number of mooring lines	3	8	8
Fairlead depth (m)	70	4	47.89
Fairlead radius from centerline (m)	5.2	28.23	27
Anchor radius from centerline (m)	853.87	423.4	27
Mooring line unstretched length (m)	902.2	473.3	151.7
Mooring line diameter (m)	0.09	0.0809	0.127
Mooring line linear density (kg/m)	77.7	130.4	116
Mooring line EA (kN)	384 243	589 000	1 500 000
Mooring line EI (kN-m^2)	38	47	300
Mooring line GJ (kN)	38	47	300

Several details about the implementation of the OC3 Hywind platform are worth noting. First, the definition of the OC3-Hywind in FAST includes the specification of additional hydrodynamic damping terms in surge, sway, heave, and yaw in order to provide for realistic damping levels that are not fully captured by the linear hydrodynamics preprocessing in WAMIT. This additional damping is included in the current work. Second, the OC3-Hywind platform has a tapered (non-constant) diameter over part of its draft. Because viscous drag effects are significant for a spar-buoy, the viscous drag calculations in FAST were altered for the OC3 studies to account for this changing diameter. This alteration was maintained in the current work. Lastly, the OC3 Hywind's mooring system features a crow-foot connection to the spar to increase

the stiffness in yaw. In the input files from NREL, the crow foot is not supported by the quasi-static mooring model, and is instead represented by an added stiffness in the yaw DOF of the hydrostatic stiffness matrix. This approach was also followed for the simulations using the ProteusDS mooring model, with the crow foot neglected in the mooring system, in order to maintain consistency when comparing with the quasi-static model.

3.2.5 Test Cases

The comparison of the dynamic and quasi-static mooring models was completed using a selection of test scenarios, or load cases (LCs), taken from the OC3 phase IV study [23]. These load cases follow the numbering scheme of the OC3 and are summarized in Table 3.2. In all cases the wind and wave directions are coincident, in the positive x direction, and all six platform DOFs are enabled.

Load case 1.4 consists of six tests, one for each platform DOF, in still water conditions with no wind. Generator and drive train flexibility DOFs are disabled. In these free-decay tests, the platform begins displaced from its equilibrium position in the respective DOF and is then released and allowed to settle to equilibrium. Table 3.3 lists the magnitude of initial displacements used, which are chosen to reflect reasonable extreme displacements for the given mooring system. Load cases 4.1 through 5.3 consist of simulations that are 20 minutes in duration. The first 10 minutes are discarded in order to avoid the inclusion of start-up transients in the analysis. Each of the test cases with irregular waves or turbulent winds is run six times with different seeds for the wind and wave randomization processes to account for the stochastic variability in the results.

Using these load cases, differences between the two mooring models can be studied for a range of different conditions: in still water, in regular and irregular waves, in steady and turbulent winds, and in situations where the wind turbine is not operating, operating below rated power, and operating at rated power.

3.3 Results

This section compares the results generated using the two different mooring line models. Emphasis is placed on the pitch DOF because it is generally the greatest contributor to wave-induced wind turbine loads, as discussed in Section 1.2. In the

Table 3.2: Load cases (LCs) considered

Test	Aero.	Wind	Waves	
LC 1.4	off	-	still	Tests the unforced transient response of the platform when perturbed in each DOF (6 tests). Aerodynamic forces are disabled.
LC 4.1	off	-	6 m, 10 s Regular	Tests the steady state response of the system in regular waves of 10 s period and 6 m height. Aerodynamic forces are disabled.
LC 4.2	off	-	6 m, 10 s JONSWAP	Tests the transient response of the system in irregular waves with a JONSWAP spectrum with 10 s peak period and 6 m significant height. Aerodynamic forces are disabled.
LC 5.1	on	8 m/s steady	6 m, 10 s Regular	Tests the steady state response of the system in regular waves and steady 8 m/s winds.
LC 5.2	on	11.4 m/s 17% TI	6 m, 10 s JONSWAP	Tests the transient response of the system in irregular waves and turbulent winds at 11.4 m/s with 17% turbulence intensity.
LC 5.3	on	18 m/s 15% TI	6 m, 10 s JONSWAP	Tests the transient response of the system in irregular waves and turbulent winds at 18 m/s with 15% turbulence intensity.

Table 3.3: Initial displacements for load case 1.4

DOF	OC3	ITI	TLP
Surge (m)	10	10	10
Sway (m)	10	10	10
Heave (m)	5	5	0.25
Roll (°)	5	5	1
Pitch (°)	5	5	1
Yaw (°)	10	10	10

figures that follow, "quasi-static" refers to results generated using FAST's default quasi-static mooring model, and "dynamic" refers to results generated using the dynamic FEM-based ProteusDS mooring model.

3.3.1 Dynamic Model Convergence and Static Equivalence

The numbers of elements in the cable discretizations for the ProteusDS mooring line model were chosen based on sensitivity studies for each of the three designs under both static and dynamic conditions. For the OC3-Hywind, a cable discretized into 20 elements was found suitable for the present work, having a static fairlead force response agreeing with a 40-element discretization to within 0.2%, and being significantly faster to compute. Similarly, a 20 element discretization was found suitable for the ITI barge and a 10 element discretization was found suitable for the TLP.

The static equivalence between the dynamic ProteusDS mooring model and the quasi-static mooring model was tested by using the first mooring line from the OC3-Hywind design (which extends in the positive x direction) and comparing the fairlead tensions once the line had come to static equilibrium for a range of surge displacements of the platform. A comparison of the resulting fairlead forces in the x (surge) and z (heave) directions is shown in Fig. 3.2. The corresponding normalized differences between the two models are shown in Fig. 3.3. The static equivalence is very good, with percentage errors below 0.6%. The variations in the relative differences (at the calculated tensions) are likely caused by the ProteusDS discretization of the mooring line-seabed interaction; when a cable element transitions on or off the sea floor there is a sudden change in the fairlead tension.

3.3.2 Free Decay Tests

Load case 1.4 is valuable in revealing the damping on the platform from platform hydrodynamics and mooring line dynamics. The tests, one for each platform DOF, are free of wind and incident waves. The decaying motions in the DOF of interest as the platform returns to equilibrium from its initial perturbed position can be processed to find the natural frequency and logarithmic decrement of the response. From the averaged logarithmic decrement, the equivalent damping ratio, ζ_d, in that DOF can be approximated. This value represents the collective damping effects of the mooring lines and the platform hydrodynamics. The results for each DOF of each design for both the quasi-static and dynamic mooring models are given in Figure 3.4.

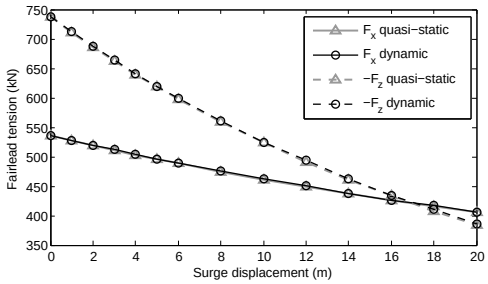

Figure 3.2: Horizontal and vertical fairlead rensions

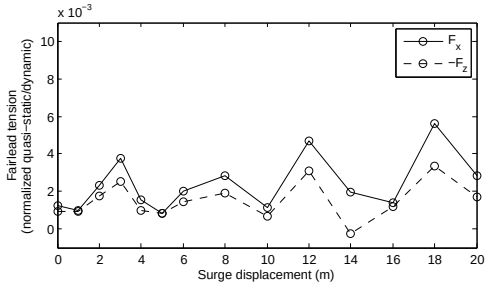

Figure 3.3: Normalized Horizontal and Vertical Fairlead Tensions

For each design the quasi-static mooring model underpredicts damping on the platform in translational DOFs, compared to the dynamic mooring model. The one exception is the TLP design in the heave direction, for which the vertical mooring lines and submerged platform provide very little damping to begin with.

The effect on rotational DOFs is highly dependent on the fairlead locations. In the OC3 design, the fairleads are located very near the platform's pitch and roll axes, making the mooring line motions and their contribution to platform damping quite small in these DOFs. Consequently, the difference in damping ratios between the two mooring models in pitch and roll is minimal. The contribution to the damping ratio

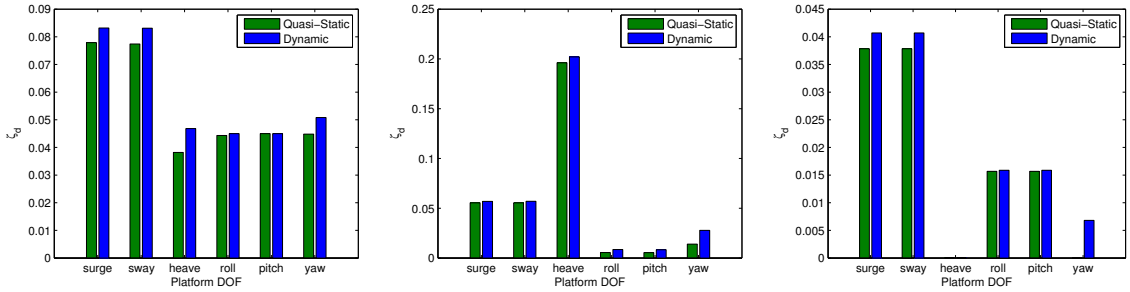

(a) OC3 (b) ITI (c) TLP

Figure 3.4: LC 1.4 - platform damping ratios

in the yaw DOF is more significant because the spar-buoy's inertia and hydrodynamic damping in the yaw DOF are small. In fact, according to the inviscid linear hydrodynamics approach there would be no hydrodynamic yaw damping from the platform; however, the OC3 design model includes supplementary linear damping terms in each DOF based on empirical results. For the ITI design, the fairleads are spaced far from the center of mass, resulting in large line motions when the platform rotates. Thus noticeable damping effects are created by the lines. Figure 3.4(b) confirms this, showing that the dynamic mooring model yields damping ratios in rotational DOFs 50% to 100% larger than those yielded by the quasi-static model. For the TLP design, the vertical orientation of the lines results in minimal drag during pitch and roll motions, which essentially just stretch and relax the lines axially. Yaw of the TLP produces significant lateral motion of the lines and hence there is significant damping when the dynamic mooring model is used. There is no yaw damping of any form in the quasi-static case because the platform is cylindrical and therefore has no hydrodynamic damping in WAMIT's inviscid analysis when rotated about its axis of symmetry.

3.3.3 Periodic Results - Platform Motions

Load case 4.1 allows comparison of the periodic steady-state response of the system without the effects of rotor aerodynamics. Load case 5.1 adds the rotor aerodynamics at a steady wind speed of 8 m/s, providing approximation of the coupled system performance in a steady-state operating condition. Across the three support structure designs being evaluated, the platform surge and pitch motions contribute in different amounts to the loads on the wind turbine. Their contribution to the acceleration of the nacelle is approximately equal in the OC3 design, pitch is dominant in the ITI design, and surge is dominant in the TLP design. Given the additional importance of pitch motions to bending loads on the blades and gyroscopic effects of the rotor, the pitch DOF is used for the comparison metric in the following results. Figures 3.5 and 3.6 show the power spectral density (PSD) plots of each platform's response in the pitch DOF for load cases 4.1 and 5.1, respectively. Agreement between models tends to be similar for both load cases.

For the OC3 design, the PSDs from the two mooring models agree to within 0.3% at the 0.1 Hz peaks corresponding to the regular 10 s period waves driving the system motions. Excitation at other frequencies is minimal (over 5 orders of magnitude less than the peak), and is a result of nonlinearities in the system.

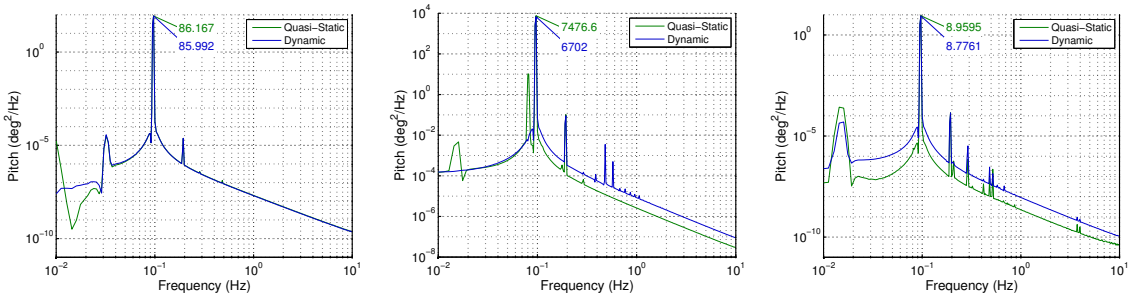

(a) OC3 (b) ITI (c) TLP

Figure 3.5: LC 4.1 - Platform Pitch PSD

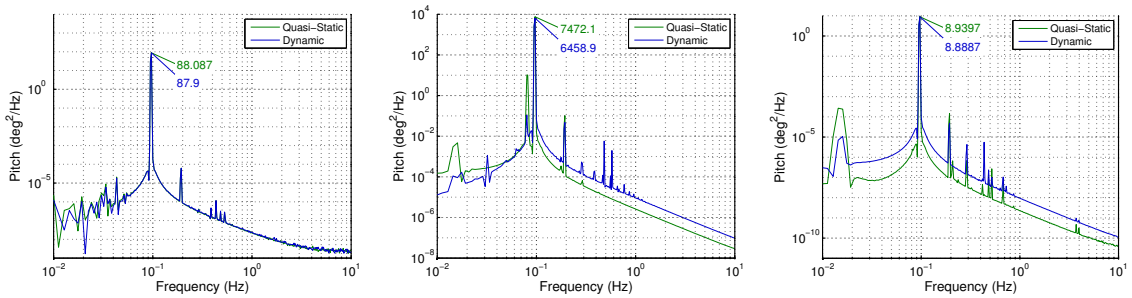

(d) OC3 (e) ITI (f) TLP

Figure 3.6: LC 5.1 - platform pitch PSD

For the ITI design, the quasi-static model overpredicts the excitation at the 0.1 Hz peak by about 10% relative to the dynamic model, but at higher frequencies underpredicts the excitation by a factor of 3. The 10% disagreement at the peak indicates the significance of dynamic mooring line effects for the ITI design.

For the TLP design, the quasi-static model overpredicts excitation at the peak by 0.3%, and underpredicts by a factor of 5 to 10 over the rest of the spectrum. This underprediction over the rest of the spectrum, which also occurred for higher frequencies with the ITI design, demonstrates how the dynamic mooring line effects lead to the introduction of higher-frequency dynamics.

3.3.4 Stochastic Results - Platform Motions

For the stochastic simulations, which have each been run six times with different realizations of the stochastic wind and wave conditions, the PSD plots are produced by averaging the PSDs of all six simulations. Load case 4.2 tests the platform response in irregular waves, without aerodynamic effects from the wind turbine. The averaged PSD plots for each design are given in Figure 3.7.

Load case 5.2 tests the platform response in irregular waves, with the effects of the wind turbine operating at its 11.4 m/s rated wind speed with 17% turbulent intensity. This case provides a comparison of the mooring line models when the system is in a normal operating state with the blade pitch and generator torque controllers in the region II-region III transition[1]. The averaged PSD plots for each design are given in Figure 3.8.

Load case 5.3 tests the platform response in irregular waves, with the effects of the wind turbine operating at rated power in 18 m/s average winds with 15% turbulent intensity, using the region III blade pitch controller. The averaged PSD plots for each design are given in Figure 3.9.

Looking at the results for the OC3 Hywind, the quasi-static mooring model agrees extremely well with the more-accurate dynamic mooring model over the majority of the excitation spectrum. The worst case of disagreement at the peaks is 2%, in load case 5.3. Two peaks are dominant in the figures: a peak at 0.03 Hz corresponding to the pitch natural frequency of the platform, and a peak at 0.1 Hz corresponding to the 10 s peak period of the wave spectrum. The presence of wind and aerodynamics has little effect on the excitation around the 0.1 Hz peak, but adds significant excitation around the 0.03 Hz peak and also produces a small peak at 0.5 Hz which is close to the 3P frequency of 0.6 Hz[2].

For the ITI Energy Barge, the dominant peak falls between 0.084 Hz and 0.1 Hz, corresponding to the 0.084 Hz pitch natural frequency of the platform and the 10 s peak period of the wave spectrum, respectively. As with the OC3 design, there is a small peak at 0.5 Hz near the 0.6 Hz 3P frequency when wind and aerodynamics are enabled. In load case 4.2, the quasi-static model significantly over-predicts the excitation across the spectrum. The peak is overpredicted by a factor of 3.4. The disagreement is more moderate in load case 5.2, with overprediction at the peak of 7%, though there is noticeable overprediction of the excitation around 0.04 Hz. Load case 5.3 shows closer agreement across the spectrum, but the peak is overpredicted by 12%.

For the MIT/NREL TLP, there is a wide region of elevated excitation spanning 0.09 Hz to 0.23 Hz (corresponding to near the peak period of the wave spectrum and

[1]"Region II" refers to wind speeds for which the wind turbine is operating under torque control, below rated power. "Region III" refers to wind speeds when the wind turbine is operating under blade pitch control, at rated power.

[2]The "3P" frequency, triple the rotation frequency of the rotor, is a common excitation frequency in three-bladed wind turbines.

the 0.23 Hz pitch natural period of the TLP, respectively). Agreement between peaks in this region is within 4% in load case 4.2 and within 1% in load cases 5.2 and 5.3 where wind and aerodynamics enabled. There are very small peaks in excitation at 0.6 Hz (the 3P frequency) and 0.7 Hz in all three load cases. The presence of this excitation in load case 4.2, which has no rotor rotation, suggests that this excitation is not related to 3P excitation. In any case, its amplitude is extremely small. There are some noticeable disagreements between the models in other regions of the spectrum, mostly in load case 4.2, but these are in regions with orders-of-magnitude lower levels of excitation.

3.3.5 Stochastic Results - Tower and Blade Loads

The blade root and tower base bending moments are two of the critical loads for floating wind turbines. The damage-equivalent loads[3] and extreme loads in these two locations as calculated using the two different mooring models are compared in Figures 3.10 and 3.11 in terms of the quasi-static result normalized by the dynamic result. Since each test is run six times with different realizations of stochastic wind and wave conditions, these figures show the average and standard deviation of the six normalized load metrics for each design in each load case. Both the normalized damage-equivalent loads and the normalized extreme loads correspond well with the varying levels of disagreement in pitch PSD between the two mooring models, for each design in each load case.

For the OC3 design, which experienced very good agreement between mooring models, the normalized values are within 1% of unity. This is fairly consistent with the results of Matha et al. [44], which show close agreement in the wind turbine loads. It is in some contrast to the results of Kallesøe and Hansen [54], which show tower fatigue loads reduced by around 5-10% with the dynamic model. Their dynamic model neglected the effects of wave kinematics and added mass on the mooring lines; this may be one of the reasons for the contrast with the present results.

The ITI design, which experienced significant pitch overprediction by the quasi-static mooring model, has normalized damage-equivalent and extreme loads significantly greater than unity. For load case 4.2, by far the worst case of disagreement, the damage-equivalent and extreme load overpredictions for both tower base bending

[3]Damage-equivalent loads are constant-amplitude loads that would result in the same amount of fatigue damage as the load time series in question. The utility MCrunch by NREL was used to calculate these quantities in the present work.

and blade root bending exceed 30%. The large standard deviation in many of the ITI results demonstrates the statistical variability inherent in simulations with these stochastic environmental conditions.

The TLP design, which experienced small pitch disagreements between mooring models, has correspondingly small deviations from unity in its normalized load values. In load case 4.2, again the worst case for disagreement, the damage-equivalent and extreme load overpredictions by the quasi-static model for blade root bending moment are 8% and 9%, respectively. This stands in contrast to the average overpredictions of the tower base damage-equivalent and extreme loads for the same load case, which are around 1%.

3.4 Discussion

The stochastic results provide some clear distinctions between different floating wind turbine designs as to whether a quasi-static mooring line model is adequate. For the OC3 Hywind, it is apparent that the choice of mooring model has very little impact on simulation results. Considering that the fairlead locations are near the pitching center of the spar and the motions of the Hywind are most significant at lower frequencies (resulting in smaller velocities and accelerations), the dynamic effects on the cables are minimal so it is not surprising that a dynamic model provides minimal benefits over a quasi-static model.

For the ITI Energy Barge, an extreme example of a buoyancy-stabilized design, the quasi-static mooring model results in over-prediction of the platform motions by a significant amount. The barge has very large motions, which are greatest near the 0.1 Hz peak wave excitation frequency, and the fairleads are located at a large radius from the platform's center of mass and center of rotation. These factors make the motions of the mooring lines quite significant and also give the mooring lines a large moment arm through which to affect the platform's pitching motion. Consequently, hydrodynamic drag and added mass of the mooring lines contribute significantly to the platform dynamics; using a quasi-static model neglects these important contributions.

For the MIT/NREL TLP, there is fairly good agreement in the platform motions, and the damage-equivalent and extreme tower base bending loads are overpredicted by the quasi-static model by no more than 1%. Greater disagreement and large standard deviations in the blade root bending loads, arising in load cases 4.2 and 5.3, suggest that the blade loads are much more sensitive than the tower loads to platform motion.

The 8% and 9% differences in the damage-equivalent and extreme blade bending loads in load case 4.2 are perhaps not representative of the true sensitivity because this load case lacks aerodynamic effects, which appear to reduce the sensitivity of blade bending loads to platform motion. All in all, the quasi-static model provides an approximation of the system dynamics with inaccuracies in tower load metrics in the order of 1% and inaccuracies in the blade load metrics of at most 10%. It seems that the larger size and tension of the mooring lines in a TLP design make the influence of their dynamic effects on the system response far from negligible.

An important observation can be made from comparing the time series data produced using the two different mooring models. Although platform pitch is generally the DOF contributing most to the increased tower and blade loads in a floating wind turbine, comparing only the predicted pitching motions runs the risk of overlooking significant differences in the higher-frequency structural dynamics of the tower and rotor blades. The first fore-aft bending natural frequency of the tower is near 0.3 Hz, and the first flapwise bending natural frequency of the rotor blades is near 0.7 Hz [24]. When the platform motion excites these higher-frequency structural modes, even small differences in the platform motion can induce differences in the blade deflection that produce large second-by-second differences in the blade bending load time series. These differences are minor for the highly-stable OC3 design, for which platform motion and tower load time series agree very well between mooring models. For the ITI and TLP designs, however, these differences are significant. This is apparent in Figures 3.12 and 3.13, which show time-series snapshots from the ITI design in load case 5.2 and the TLP design in load case 5.3. Progressively-increasing levels of disagreement are visible in the time series as one moves up the structure, from platform pitch to tower base bending moment to blade deflection and blade bending moment. Fairly large differences are visible in certain peaks in the blade bending load time series (see, for example, Figure 3.12d at $t = 853\,\mathrm{s}$ and Figure 3.13d at $t = 932\,\mathrm{s}$).

Though the quasi-static mooring model is more likely to overpredict a given peak, there are cases where it underpredicts a peak. In a few of the simulations, the extreme blade bending load produced using one mooring model occurs at a completely different time than the one produced using the other model. Consequently, large time series differences do not necessarily translate into large damage-equivalent or extreme load differences. Consider, for example, how the maximum loads in Figure 3.13d (at $t = 892\,\mathrm{s}$) disagree by about 15%, yet the average normalized extreme load calculated for the full set of TLP load case 5.3 simulations in Figure 3.10d indicates disagreement

of only 3%. Looking more broadly, the instantaneous blade bending load differences are as large as 50% of the mean load for the ITI design and 30% of the mean load for the TLP design, yet these time series differences manifest in the calculated damage-equivalent and extreme loads as differences of at most 30% for the ITI design and at most 10% for the TLP design. As such, comparing damage-equivalent and extreme load metrics can provide a floor to the measure of disagreement between mooring models; comparing them will not capture the full extent of time series differences.

Regardless of the level of disagreement in the calculated load metrics, the differences in Figures 3.12d and 3.13d demonstrate that high-accuracy mooring modelling is crucial for accurate prediction of the blade load time series. The lack of well-scaled experiments on floating wind turbines prevents the confirmation of these findings with experimental results.

3.5 Conclusions

The dynamic FEM-based mooring line model ProteusDS was coupled to the floating wind turbine simulator FAST to provide a means of comparing quasi-static and dynamic mooring models. The static equivalence of the ProteusDS model and FAST's built-in quasi-static mooring model is very good. Three floating wind turbine designs, one from each stability class, were tested using the two mooring models in a number of load cases, including free decay tests, periodic steady-state operating conditions, and stochastic operating conditions.

From the free decay tests, there is clear evidence that the damping from the mooring lines makes up a significant portion of the overall platform damping in some cases. Which DOFs are most effected, and the extent of the contribution, depends on the mooring line orientations and the fairlead locations, so it varies considerably between support structure designs.

From the results of the stochastic load case simulations, a number of support structure-specific conclusions can be made:

- Quasi-static mooring models are well-suited for slack-moored designs that have natural periods well below the peak wave periods, small motions, and fairleads located close to the platform's center of mass (eg. spar-buoys).

- Quasi-static mooring models are not suited to slack-moored designs that have

natural periods near or above the peak wave periods, large motions, and fairleads located far from the platform's center of mass (eg. barges).

- The suitability of quasi-static mooring models for slack-moored designs that fall in between the two above extremes in terms of natural periods, motion amplitudes, and fairlead locations should be evaluated on a case-by-case basis using a comparison study similar to this one.

- For taught-moored designs, quasi-static mooring models can provide an approximation of the system dynamics but cannot provide high accuracy in turbine load prediction.

Regardless of support structure, it has been shown that small inaccuracies in the platform motion time series introduced by a quasi-static mooring model can cause much larger inaccuracies in the time series of the higher-frequency rotor blade dynamics. Large second-by-second inaccuracies in rotor blade bending load do not *necessarily* translate into large differences in the corresponding damage-equivalent and extreme loads calculated over multiple stochastic simulations. Consequently, differences in damage-equivalent and extreme load metrics should be considered a floor to the measure of inaccuracy caused by a quasi-static mooring model; greater instantaneous inaccuracies will be present in the time series response. These time-series inaccuracies should be given careful consideration if using other design metrics that may be more sensitive to time series differences.

To conclude, this analysis has shown that the adequacy of quasi-static mooring models depends heavily on the support structure configuration being simulated. For a given support structure, accurate mooring modelling is of elevated importance for predicting the damage-equivalent and extreme bending loads on the rotor blades and is crucial for accurately predicting the blade bending load time series.

The work discussed in this chapter was a look into one of the lesser-explored modelling accuracy questions in floating wind turbines. It provides a higher-fidelity mooring model capability to the current industry standard in floating wind turbine simulation, provides some valuable general conclusions about the suitability of quasi-static mooring models in some cases and not others, and also demonstrates a systematic approach by which the mooring model fidelity required for a given floating wind turbine design can be determined. The work relied on relatively high-fidelity

time-domain simulation. The work in the following chapters takes a step back from modelling issues and turns its attention to the question of support structure design.

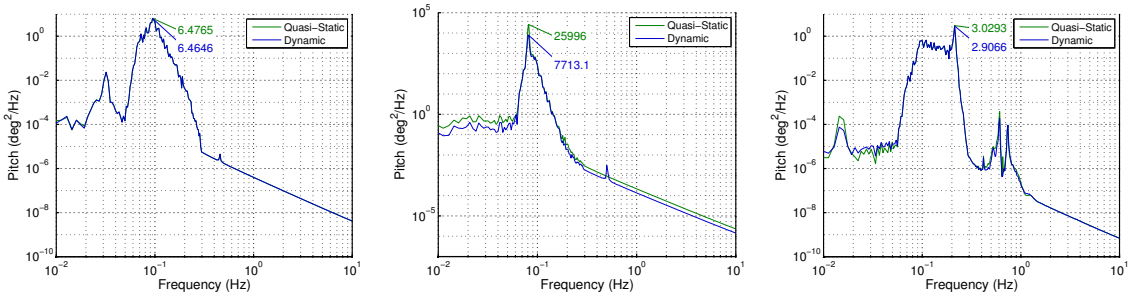

(a) OC3 (b) ITI (c) TLP

Figure 3.7: LC 4.2 - platform pitch PSD

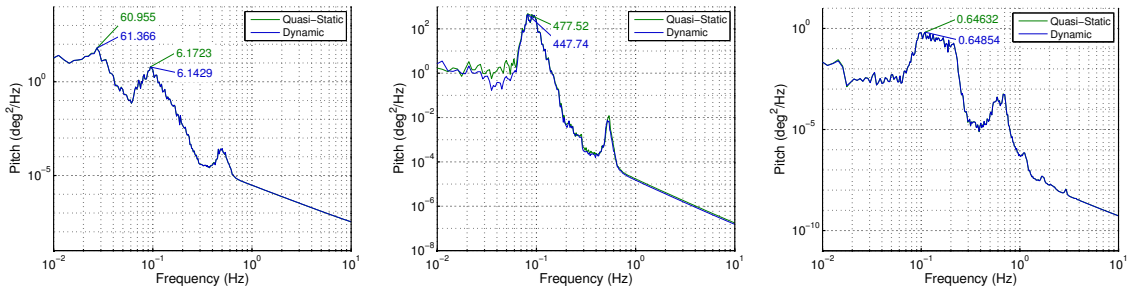

(d) OC3 (e) ITI (f) TLP

Figure 3.8: LC 5.2 - platform pitch PSD

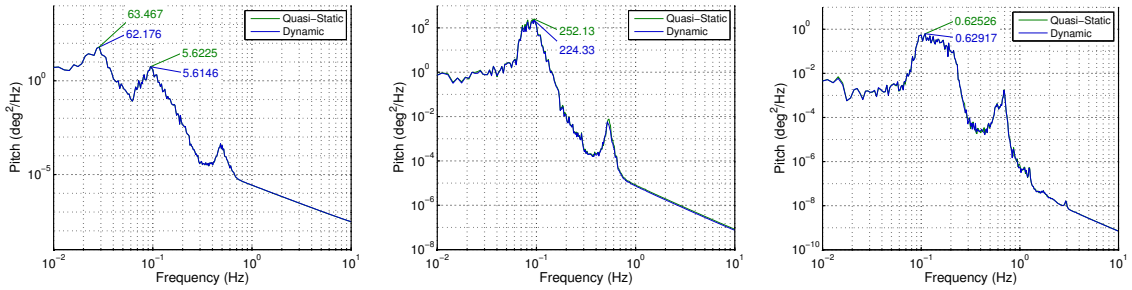

(a) OC3 (b) ITI (c) TLP

Figure 3.9: LC 5.3 - platform pitch PSD

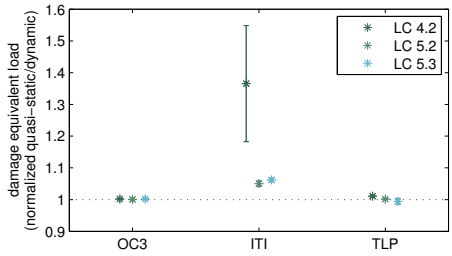

(a) Tower base bending moment

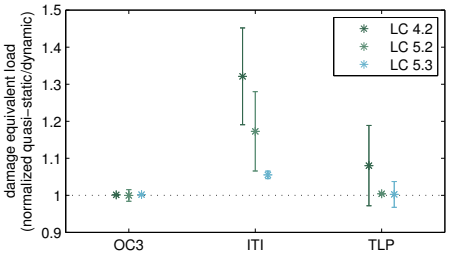

(b) Blade root flapwise bending moment

Figure 3.10: Damage equivalent loads

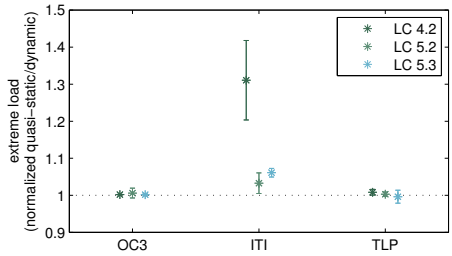

(c) Tower base bending moment

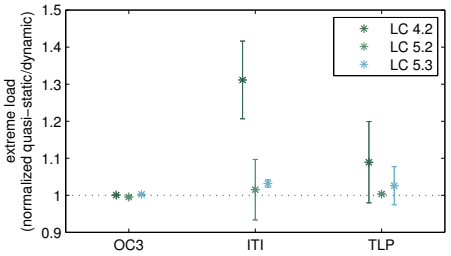

(d) Blade root flapwise bending moment

Figure 3.11: Extreme loads

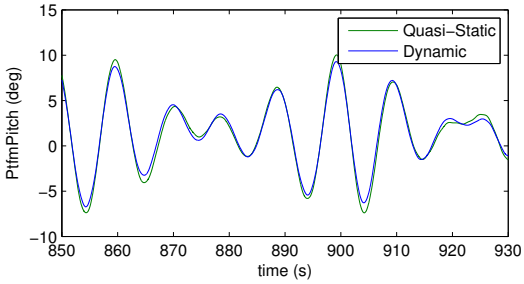

(a) Platform pitch

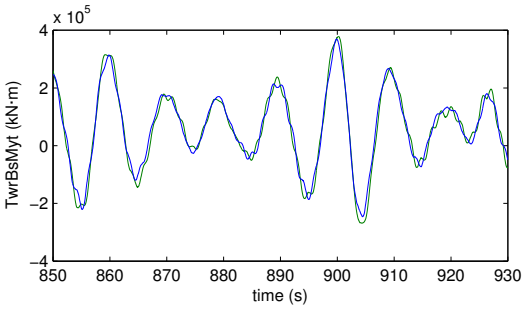

(b) Tower base bending moment

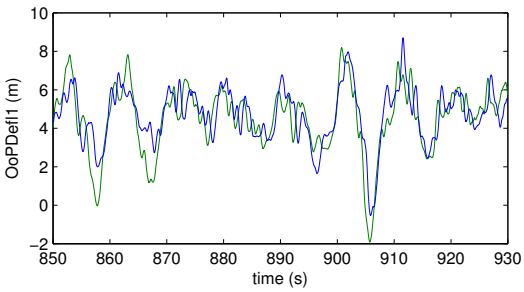

(c) Blade deflection in out-of-plane mode 1

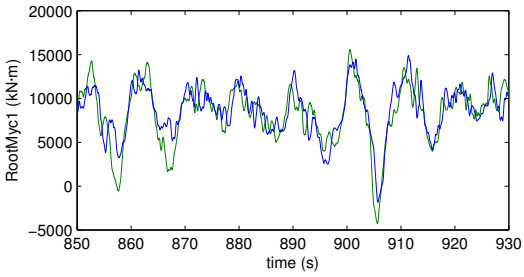

(d) Blade root flapwise bending Moment

Figure 3.12: Selected time series of ITI Energy Barge in LC 5.2

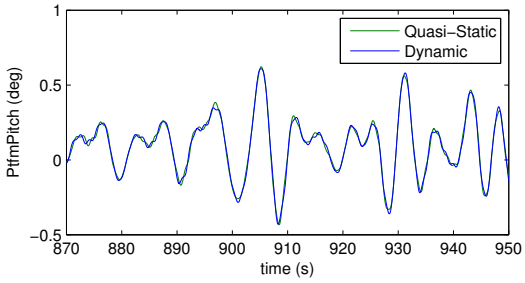

(a) platform pitch

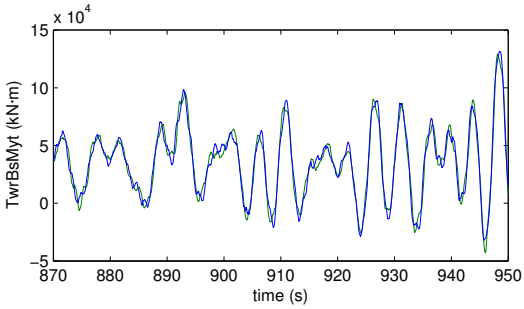

(b) Tower base bending moment

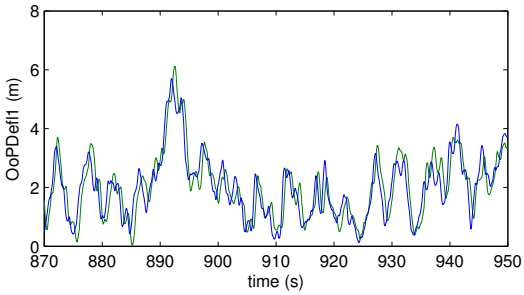

(c) Blade deflection in out-of-plane mode 1

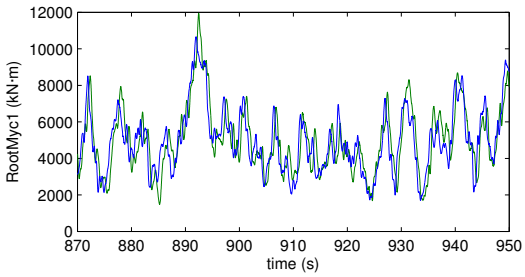

(d) Blade root flapwise bending Moment

Figure 3.13: Selected time series of MIT/NREL TLP in LC 5.3

Chapter 4

Hydrodynamics-Based Platform Optimization – A Basis Function Approach

4.1 Introduction

The variety of designs proposed for floating wind turbines and the lack of convergence to a single optimal configuration is evidence to the intractability of the support structure design problem. As Chapter 1 discussed, the support structure design problem is characterized by significant technical challenges, conflicting design objectives, and innumerable configuration possibilities. While many design alternatives are being explored and compared, most of the existing designs are based on concepts from the offshore oil and gas industry, and attempts to explore the design space systematically have so far been quite limited. This means that large regions of the design space have so far been ignored. Meanwhile, the lack of conclusive comparison between designs means that R&D attention is being divided among a number of alternative platform configurations.

A global optimization tool capable of operating on the complex floating wind turbine design problem could hold a lot of value in both exploring unexamined parts of the design space and confirming the optimality of existing design efforts. This could help to guide and unify R&D efforts and ultimately accelerate the convergence of the technology.

This chapter describes the creation, execution, and evaluation of a so-called "hydrodynamics-

based" optimization approach. Such an approach operates on the hydrodynamics characteristics rather than the underlying geometry of candidate designs, bypassing the computationally-demanding process of calculating hydrodynamic coefficients. It was hoped that this approach could explore the design space more intuitively, efficiently and broadly than previous efforts. If successful, such an approach could simplify and clarify the otherwise complex and obfuscated support structure design problem, and lead to the identification of key platform hydrodynamic characteristics that maximize performance.

The method cannot yield finalized designs, or be entirely faithful to the physics involved, however the exercise provides some valuable insights about the optimization problem and can serve as a starting point and motivator for further pursuits of hydrodynamics-based optimization.

4.1.1 Conventional Geometry-Based Design Space Exploration

There is still no consensus as to which stability class or platform configuration holds the most promise. Studies have been done in the past both comparing and optimizing floating platform designs. Comparison efforts of specific designs can afford to use computationally-intensive time-domain simulation tools to provide a detailed and reliable comparison of leading designs from each of the three stability classes. See for example the work of Jonkman and Matha [50] and Robertson and Jonkman [22]. While these studies are an excellent later-stage tool for identifying the best design, they lack the ability to explore the design space for new design concepts. Optimization efforts and parameter studies, which are capable of design space exploration, tend to rely on lower-fidelity computationally-efficient frequency domain modelling techniques.

The conventional method of finding an optimal floating platform shape is by parametric optimization, in which the geometry of the platform is represented by a number of parameters that become decision variables. Parametric approaches are attractive because they reduce the design variables to a manageable number so that it is possible to analyze and conclusively compare design alternatives. While this works well for determining the optimal dimensions of a single design concept, it is difficult to create a parameterization that can describe different configurations (eg. both a spar-buoy design and a three-column semisubmersible). Consequently, parameterizing such a complex design problem tends to artificially constrain the design space and

limit design creativity.

The best example in the literature of a parameter study that considers a broad range of platform configurations is that of Tracy at MIT [51]. The parameterization uses a cylindrical platform of variable dimensions and mooring lines of variable tension and angle – thereby including designs from each stability class, from TLPs to spar-buoys to cylindrical barges – and frequency-domain modelling to find Pareto-optimal support platform configurations. Unfortunately, this parameterization is still limited to single-cylinder configurations; multi-cylinder configurations – an important part of the design space given the recent trends toward semi-submersible platforms – are excluded.

4.1.2 Hydrodynamics-Based Optimization

The idea of "hydrodynamics-based" optimization is to represent the design space in terms of hydrodynamic performance characteristics rather than geometric characteristics. The motivation behind this idea is to represent the design space in a way that is more intuitively related to the performance characteristics that a designer needs to be mindful of. By avoiding the assumptions of a geometric parameterization, a hydrodynamics-based approach may be able to explore the design space more widely. By avoiding in-the-loop calculation of hydrodynamic properties from platform geometry (using eg. WAMIT), a hydrodynamics-based approach will be significantly faster than conventional optimization approaches. The strategy is to look for optimal support platform performance characteristics without making a priori assumptions about the platform geometry that would limit the design space. In order to achieve this, the geometric decision variables that describe the geometry of the platform need to be replaced with hydrodynamic performance-related decision variables that describe the hydrodynamic characteristics of the platform (such as hydrostatic stiffnesses and wave excitation coefficients). The challenge is to apply constraints to these characteristics so that they are realistic, without simultaneously making assumptions that over-constrain the geometry of the platform. If this can be achieved, then we will have a powerful new way of exploring the design space – a process that will yield optimal sets of platform hydrodynamic characteristics that can then be used as performance targets for more detailed support platform design work.

The most direct approach toward "hydrodynamics-based" optimization would be to treat each of the hydrodynamic parameters of the platform - added mass, damp-

ing, stiffness, and wave-excitation - as decision variables. The first problem with this is that these are all matrices or vectors, and most of them are frequency dependent, rendering the problem domain impossibly large. The second problem is that without imposing constraints on, and between, these variables, the majority of the design space would be completely unphysical and not represent the the real design configurations. But to follow the conventional approach of generating hydrodynamic coefficients from specific platform geometries would mean losing the generality and insight that hydrodynamics-based optimization aims to achieve. The task, then, is to find a middle ground – to identify a small number of generic platform properties that can be used as decision variables, that are sufficiently detailed to provide plausible estimation of hydrodynamic coefficients, and that are still general enough to represent the full floating platform design space.

4.2 Basis Function Optimization Approach

A "basis function" approach is one possibility for hydrodynamics-based optimization that offers a simple means of representing the physical constraints. The idea is to use a collection of unique geometrically-defined platform designs as "basis designs" whose hydrodynamic performance coefficients can be linearly combined to approximate the characteristics of any platform in the design space. This approach is analogous to basis functions defining a function space or basis vectors defining a vector space.

The approach relies on a frequency-domain model and the theory of linear hydrodynamics, which represents the hydrodynamics using linear frequency-dependent coefficients. To model the design space, these frequency-dependent coefficients for each basis design are superimposed to yield coefficients for new designs. Non-frequency-dependent platform mass and hydrostatic restoring matrices are also linearly combined. The coefficients of the combination, x_i, are the design variables; they are constrained to the range $[0, 1]$ and together sum to 1. The case of $x_i = 1$ dictates that basis design i fully constitutes the platform configuration; the case of $x_i = 0$ dictates that basis design i has no bearing on the platform configuration.

To illustrate, a candidate platform hydrodynamic damping matrix is calculated as

$$\mathbf{B}_o(\omega) = x_1\mathbf{B}_1(\omega) + x_2\mathbf{B}_2(\omega) + ... + x_n\mathbf{B}_n(\omega), \tag{4.1}$$

and a candidate hydrostatic stiffness matrix is calculated as

$$\mathbf{C}_o = x_1\mathbf{C}_1 + x_2\mathbf{C}_2 + ... + x_n\mathbf{C}_n, \tag{4.2}$$

where the subscripts refer to the index of the basis design. The matrices $\mathbf{B}_i$ and $\mathbf{C}_i$ are six-by-six, corresponding to the six platform DOFs.

This linear combination approach provides a straightforward way of approximating the complex interdependencies and constraints between the different hydrodynamic characteristics of a floating structure. The resulting superposition of frequency-dependent coefficients can be viewed as an extension of the linear hydrodynamics assumptions, because linear hydrodynamics relies on superposition. The question is whether the approximation is reasonable. If the frequency-dependent wave excitation curves of two different platform geometries are superimposed, will the resulting curve be physically-plausible? Given two arbitrary platform geometries, does a third geometry with intermediate hydrodynamic characteristics necessarily exist? Some example cases are considered at the end of this chapter to make a start at answering these questions.

4.2.1 Basis Platform Designs

Six platform geometries were selected to serve as basis designs. The starting point is the most simple (albeit non-ideal) geometry: a cylinder of unity aspect ratio. From there, the geometry is altered along a number of paths – elongation, splitting into multiple cylinders, etc. – to achieve different platform configurations. These are detailed in Figure 4.2 and specifications are provided in Table 4.1.

A mass model was created to prescribe the mass, center of mass, moments of inertia, and amount of ballast for each basis design based on its geometry. The model assumes a 0.18 m thick steel skin over the surface area of the platform and concrete ballast that can fill up the volume within the shell from the bottom up. The steel thickness was selected by analyzing a number of published floating turbine design specifications [25, 55, 50].

Suitable dimensions for each basis design were found by parametrically varying the dimensions and tracking the changing displaced volume, platform weight, and hydrostatic stability coefficients. Platform dimensions and use of ballast were selected so as to meet the required net buoyancy for the mooring system and certain basic

stability requirements. The spar, being the least capable of active ballast[1], was sized to keep the static pitch angle below 5 degrees at all times. The other designs were then sized to have the same displaced volume as the spar. This is a way of maintaining approximate cost equivalence between platforms, because displaced volume is proportional to the sum of the structural weight and the vertical tension in the mooring lines, which are both major cost drivers [51]. As an example, the sizing algorithm used to produce the Ring design is given in Figure 4.1. The approach could be extended with basis platforms with different displacements to allow for optimization of platform size.

- given desired volume and net buoyancy
- for incrementally increasing ring outer radius
  - calculate draft as proportion of outer radius
  - for incrementally increasing ring radial thickness
    - calculate volume, water plane area, surface area, etc.
    - apply mass model to surface area to calculate mass properties
    - calculate hydrostatic stiffness
    - if volume exceeds desired volume, break
  - if net buoyancy meets or falls under required net buoyancy, break
- calculate remaining mass properties
- discretize platform surface and create WAMIT .gdf file
- run WAMIT to obtain platform's hydrodynamic coefficients

Figure 4.1: Sizing algorithm for Ring basis platform design

4.3 Modeling and Evaluation Methodology

A linear frequency-domain model was used to evaluate points in the design space, accounting for the dynamics of the wind turbine, mooring system, and floating platform. Loads from steady winds and regular (i.e. monochromatic) waves are included. The

[1]Active ballast refers to a movable ballast, often water, which can be shifted to the side of the platform to counter the static overturning moment caused by the thrust on the wind turbine.

Table 4.1: Basis platform specifications

Basis Design Index	i	1	2	3	4	5	6
		Specifications for Slack Catenary Mooring					
		Spar	Cylinder	Barge	Semisub	Ring	Sub
water plane area	$A_{wp}(m^2)$	63.6	357.5	2275	211.1	908.3	33.2
water plane moment	$I_{wp}(m^4)$	321.5	10171	411855	286204	446734	87.6
displaced volume	$\forall(m^3)$	7627	7627	7627	7627	7627	7627
platform mass	$m(t)$	7229	7230	7230	7229	7230	7231
center of mass depth	$h_{CM}(m)$	76.2	15.4	1.68	16.1	4.2	24.3
draft	h	120	21.34	3.35	36.12	8.4	31.5
ballast mass	$m_b(t)$	2258	4200	0	2078	0	3990
		Specifications for Tension Leg Mooring					
		Spar	Cylinder	Barge	Semisub	Ring	Sub
water plane area	$A_{wp}(m^2)$	127.8	569.5	3840	198	1385	33.2
water plane moment	$I_{wp}(m^4)$	1300	25811	1173746	268398	1204896	87.6
displaced volume	$\forall(m^3)$	15335	15335	15335	15335	15335	15335
platform mass	$m(t)$	12092	12093	12093	12093	12093	12094
center of mass depth	$h_{CM}(m)$	81	19.9	2	41	5	31.7
draft	h	120	26.9	4	77.4	11.1	40
ballast mass	$m_b(t)$	4935	7265	0	2078	0	6974

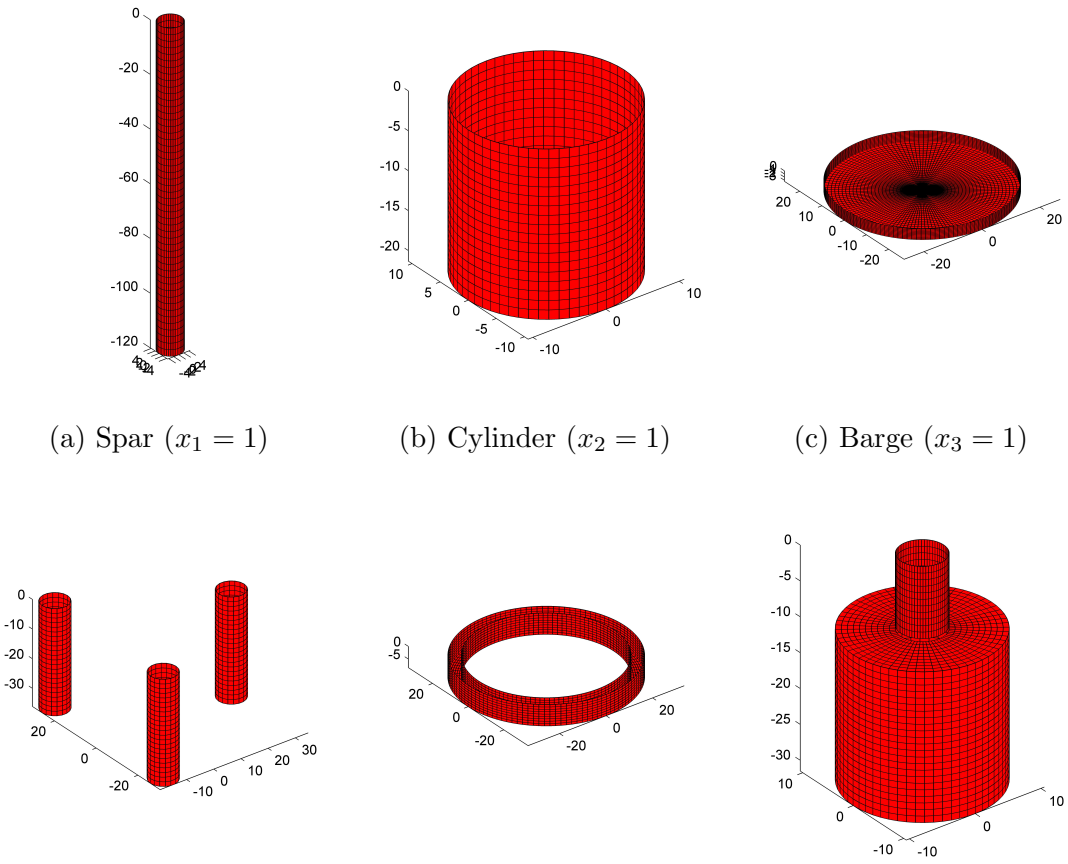

(a) Spar ($x_1 = 1$) (b) Cylinder ($x_2 = 1$) (c) Barge ($x_3 = 1$)

(d) Semisub ($x_4 = 1$) (e) Ring ($x_5 = 1$) (f) Sub ($x_6 = 1$)

Figure 4.2: Basis platform geometries for slack catenary mooring

degrees of freedom considered by the frequency-domain model are the six rigid-body modes of the platform.

A consequence of the steady-state assumption of the frequency domain approach is that the platform motions will be at the same frequency as the incident waves, and the incident waves must be regular (i.e. monochromatic). These assumptions limit the applicability of the approach because they ignore irregular wave conditions and cannot model the transient response of the system. However, the approach is computationally efficient, and the responses at different frequencies can be superimposed according to a wave spectrum to approximate the system behaviour in irregular sea states.

As described in Section 2.3.4, the generic equation of motion of a floating body (2.8) can be simplified using the steady-state and harmonic-motion assumptions of a frequency-domain model into:

$$\{-\omega^2[M + A(\omega)] + i\omega B(\omega) + C\}\boldsymbol{\Xi} = A_{wave}X(\omega) \tag{4.3}$$

The complex response, $\boldsymbol{\Xi}$, can then be solved for at any frequency if the coefficient matrices are known. The frequency-dependent response for unit amplitude waves, in terms of DOF amplitudes and phases, is commonly referred to as a response amplitude operator (RAO).

4.3.1 Hydrodynamic Loads

For the floating platform dynamics, the theory of linear hydrodynamics (as discussed in Section 2.3.3) is used, which provides the frequency-dependent coefficients that can directly be applied to equation (4.3). The added mass, **A**, wave-radiation damping, **B**, hydrostatic stiffness, **C**, and wave excitation coefficients, **X**, are calculated using the panel method potential flow solver WAMIT. Coefficients were generated for each of the basis designs considered, over a frequency range of 0.05 to 5 rad/s in 0.05 rad/s increments.

4.3.2 Wind Turbine Loads

The aero-elastic effects of the wind turbine can be added to the frequency domain model if they are linearized. This is the approach used in the Pareto-optimization work of Sclavounos et al. [16] and the approach used here. The wind turbine used is the NREL offshore 5 MW baseline turbine, as described in Section 2.1. In keeping with the linearized approach, the wind turbine mass, damping, and stiffness coefficients are constants and simply add on to the varying support platform coefficients in the equation of motion of the overall system.

Using the linearization capability of FAST, linear mass, damping, and stiffness coefficients were obtained for the wind turbine for each wind speed condition under consideration. There is potential for inconsistency in that the linearized properties depend on the static pitch angle of the platform, which will be different for each platform design at each wind speed. To account for this would require linearization within the optimization loop. For practicality, the approximate approach of linearizing about a zero pitch angle was used; the zero static pitch angle is assumed throughout the optimization.

4.3.3 Mooring System Loads

Two mooring systems have been used in this work: a 3-line slack catenary system and an 8-line vertical tension leg system, with specifications for the latter taken from the MIT/NREL TLP design specified by Jonkman and Matha [50]. Both systems use a 200 m water depth. Details are provided in Table 4.2 below, using a cylindrical coordinate system with z being positive-up measured from the waterline and r being the radius measured horizontally from the center of the platform. Of importance to the platform design is the downward force from the mooring lines, F_z, which determines the net buoyancy required.

The two systems are quite different. The slack catenary design functions to provide station keeping to the floating platform, but otherwise provide minimal stiffness to transient or periodic motions. The tension leg design functions as station keeping and also a significant source of pitch, roll, and heave stability; its taut lines provide very high stiffnesses in these DOFs.

Table 4.2: Mooring system specifications

	Symbol	Slack Catenary	Tension Leg
number of lines	n_{lines}	3	8
fairlead depth	z_{fair} (m)	-10	-47.89
fairlead radius	r_{fair} (m)	10	27
anchor radius	r_{anch} (m)	600	27
unstretched line length	$L_{unstr.}$ (m)	630	151.73
vertical load from lines	F_z (kN)	1 252	30 670

Similarly to the aerodynamic loads, the loads from the mooring lines can be linearized and added to the frequency domain model. For greatest accuracy, the mooring line reaction forces should be linearized about the steady-state operating position of the platform rather than its initial unloaded position [51]. The steady-state operating position depends on the wind speed (wind turbine thrust load), the platform hydrostatics (for static pitch angle), and the mooring system. An assumption of zero-pitch was already made to simplify the wind turbine linearization, and the mooring model was further simplified by linearizing about the undisplaced (zero-wind) equilibrium position of the platform.

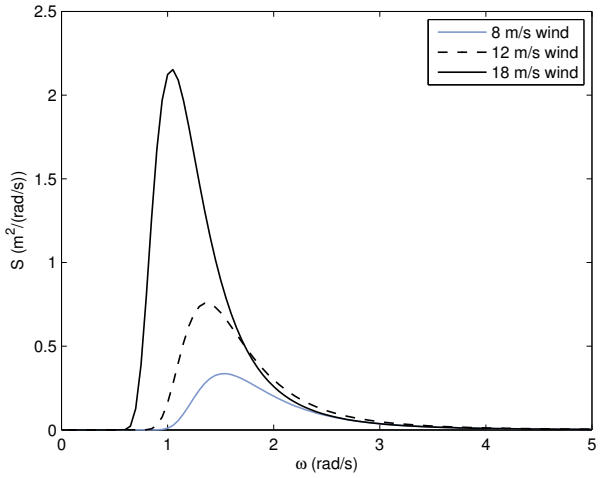

Figure 4.3: Power spectral density plots of the sea states corresponding to each wind speed

4.3.4 Environmental Conditions

The linear treatment of the system requires that steady winds and regular waves be used in the response evaluation. However, following the assumptions of linearity, the wave frequency-dependent response can be combined with a wave spectrum to predict the response under irregular wave excitation.

Three sets of steady wind speeds and irregular sea states are used in the dynamics evaluation. A JONSWAP wave spectrum is used with significant wave heights (H_s) of 2.4, 3.4, and 5 m, and peak periods (T_p) of 4.1, 4.6, and 5.1 s, respectively. The power spectral density, $S(\omega)$, of the JONSWAP spectrum is defined by

$$S(\omega) = \frac{1}{2\pi}\frac{5}{16}H_s^2 T_p \left(\frac{\omega T_p}{2\pi}\right)^{-5} exp\left[-\frac{5}{4}\left(\frac{\omega T_p}{2\pi}\right)^{-4}\right]. \tag{4.4}$$

The power spectral density of the three sea states is shown in Figure 4.3. The respective wind speeds used are 8 m/s, 12 m/s, and 18 m/s.

4.3.5 Objective Function

The critical wind turbine loads most affected by support platform motion are blade root flapwise bending moment and tower root fore-aft bending moment. Both are closely linked to platform surge and pitch motions. Flapwise bending at the blade roots can potentially be the most critical point of failure for a floating wind turbine

because the higher natural frequency and flexibility of the blades in the flapwise direction make them easily excited by nacelle accelerations. Nacelle acceleration is therefore used for the objective function. With the small angle assumptions inherent in the linear analysis, the response amplitude operator for nacelle fore-aft acceleration is [51]

$$RAO_{a\,nac.}(\omega) = -\omega^2[RAO_1(\omega) + z_{nac.}RAO_5(\omega)], \tag{4.5}$$

where $z_{nac.}$ is the nacelle height and the numerical subscripts denote the platform degrees of freedom – 1 being surge and 5 being pitch. Applying a wave power spectrum, using the theory described in Section 2.3.4, the root-mean-square nacelle acceleration in the specified wave conditions can be calculated as:

$$\sigma^2_{a\,nac.} = \int_0^\infty |RAO_{a\,nac.}(\omega)|^2 S(\omega) d\omega \tag{4.6}$$

This value, once summed across the three different sea states used, is the objective function used here. It follows the form of objective function for offshore structure optimization recommended by Clauss and Birk [37], using the concept of "significant" amplitude – analogous to the measure of significant wave height for irregular waves. The optimization was set up in Matlab using the fminsearch function and a grid of starting points to check for multiple local minima.

4.4 Optimal Platform Solutions

Separate sets of basis geometries were designed for each mooring system to account for their different net buoyancy requirements, and a separate platform optimization was done for each. The results are summarized in Table 4.3.

4.4.1 Result for Slack Catenary Mooring

With the slack catenary mooring system, the design space had numerous local minima, as evidenced by the multiple optimization solutions found from different initial starting points. Two platform configurations were notable due to their stark geometric differences and low objective function values. The first, "Design 1," is mostly Spar (52%) complemented by Semisub (22%) and Sub (23%). The other, "Design 2," is dominantly Semisub (83%) with a small amount of Semisub (14%). These op-

timal combinations are illustrated in Figure 4.4 using transparency to represent the proportion of each basis platform in the solution. Full details of the design vectors and synthesized platform characteristics are given in Table 4.3. Plots of the wave excitation, added mass and damping, and RAO for each design in the pitch DOF are provided in Figures 4.5 to 4.8. Although all of the DOFs are considered in the analysis, pitch is the most critical DOF for the objective function so only pitch is shown.

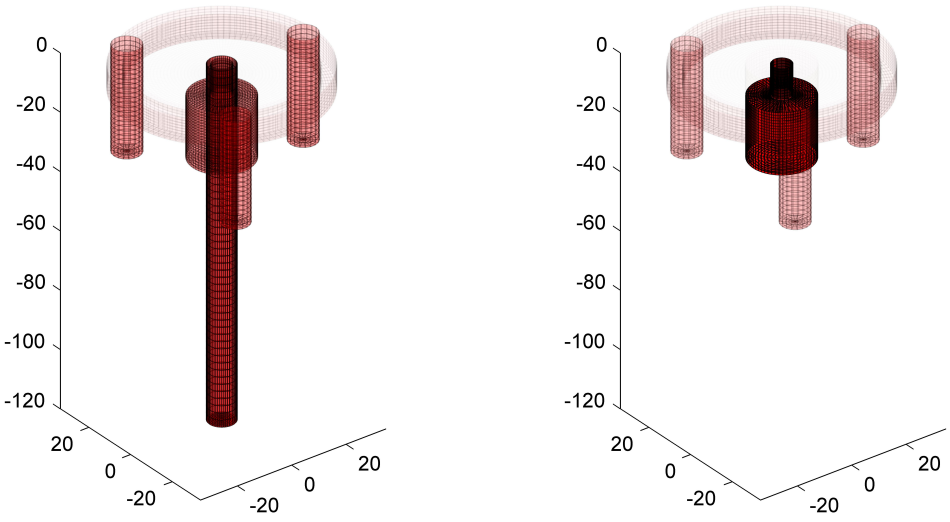

(a) Design 1 platform composition (b) Design 2 platform composition

Figure 4.4: Results for slack catenary mooring

The trend apparent the composition of the two minima is the combination of a high volume, low water plane area geometry (the Spar or Sub) with a highly-distributed water plane area geometry (the Semisub). The Semisub geometry provides a high level of hydrostatic stability but also experiences significant wave excitation inside the bandwidth of the incident wave spectra. Adding the Sub or the Spar has the effect of reducing the water plane area and thus lowering the natural frequencies of the platform away from the active wave frequencies.

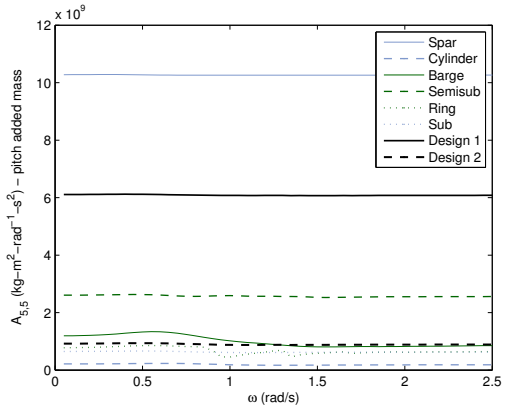

Figure 4.5: Pitch added mass for catenary-moored platforms

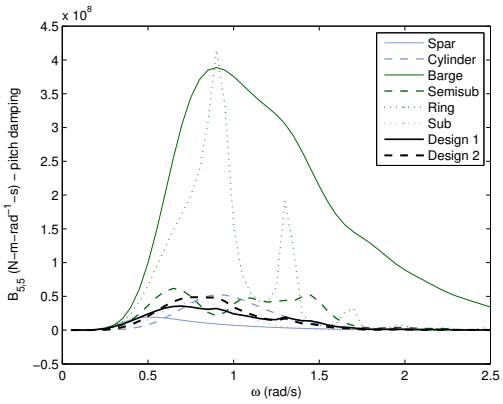

Figure 4.6: Pitch damping for catenary-moored platforms

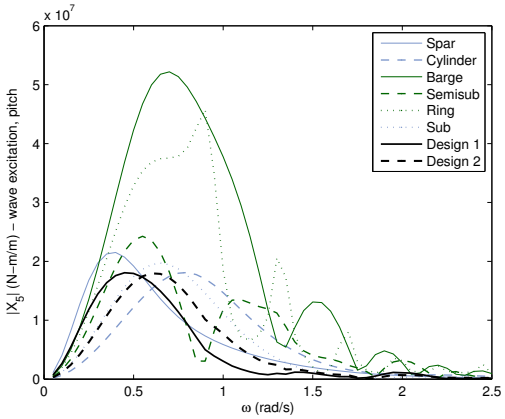

Figure 4.7: Pitch wave excitation for catenary-moored platforms

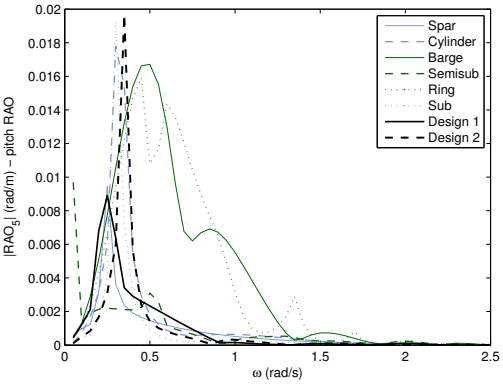

Figure 4.8: Pitch RAO for catenary-moored platforms

4.4.2 Result for Tension Leg Mooring

With the tension leg mooring system, two strong local minima, similar to those with the catenary mooring, were again observed. "Design 3" is composed of 58% Spar and 39% Semisub, and "Design 4" is composed of 83% Sub and 14% Semisub. The optimal combinations of basis platforms are illustrated in Figure 4.9. Details are provided in Table 4.3. Plots of the wave excitation, added mass and damping, and RAO for each design in pitch are provided in Figures 4.10 to 4.13. Looking at the wave excitation coefficients in Figure 4.12, it is apparent that having the excitation peak at the lowest-possible frequency is a major driver of the basis design combination; the Spar, Semisub, and Sub have the lowest-frequency peaks. The surplus of static stability provided by the tension leg mooring system does not appear to reduce the demand

Table 4.3: Optimization results

		Catenary Mooring		Tension Leg Mooring	
	Symbol	Design 1	Design 2	Design 3	Design 4
Spar	x_1	0.5223	0	0.5778	0.0015
Cylinder	x_2	0	0.0001	0.0041	0.0001
Barge	x_3	0.002	0.002	0	0
Semisub	x_4	0.2171	0.1359	0.3903	0.1392
Ring	x_5	0.0323	0.0338	0.0278	0.0245
Sub	x_6	0.2262	0.8281	0	0.8346
objective function	σ (m^2/s^4)	0.0048	0.0044	0.0053	0.0088
water plane area	A_{wp} (m^2)	120.4	91.5	192.0	89.5
water plane moment	I_{wp} (m^4)	77575	54892	139109	66958
displaced volume	$\forall$ (m^3)	7627	7627	15335	15335
platform mass	m (tonnes)	7230	7230	12092	12093

for the well-distributed water plane area of the Semisub. Perhaps the additional pitch damping provided by the Semisub platform is beneficial in this case, to damp the stiff response from the taut mooring lines.

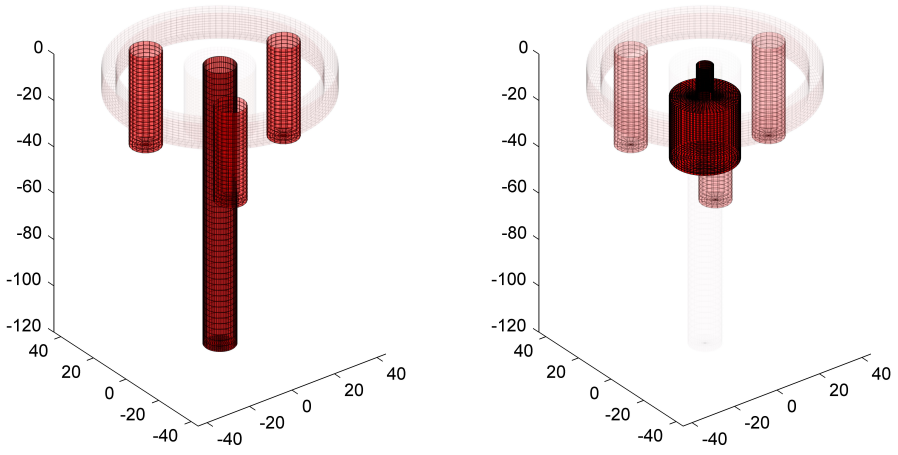

(a) Design 3 platform composition (b) Design 4 platform composition

Figure 4.9: Results For Tension Leg Mooring

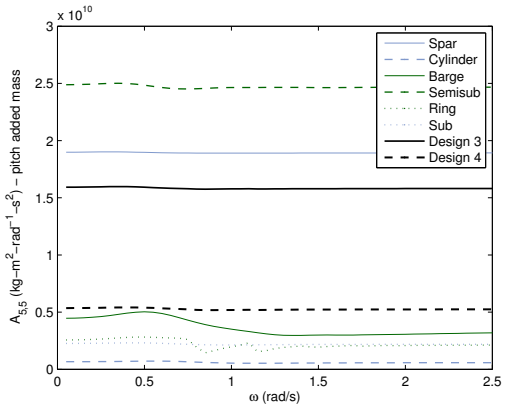

Figure 4.10: Pitch added mass for tension-leg-moored platforms

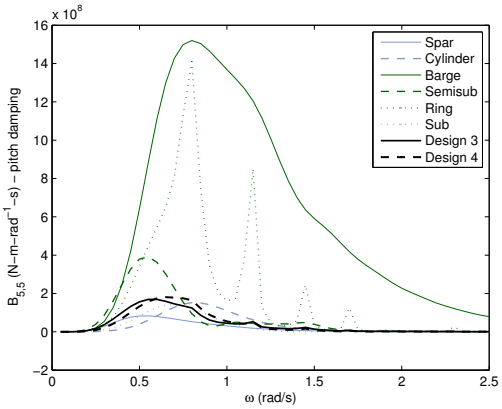

Figure 4.11: Pitch damping for tension-leg-moored platforms

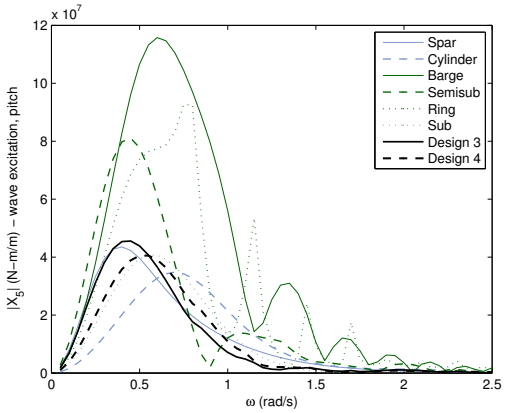

Figure 4.12: Pitch wave excitation for tension-leg-moored platforms

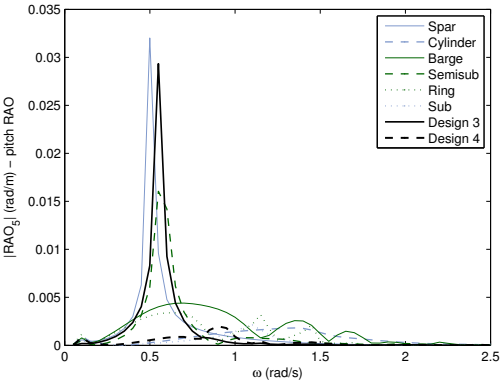

Figure 4.13: Pitch RAO for tension-leg-moored platforms

4.5 Discussion of Physical Interpretations

Weaknesses in the conceived basis function approach begin to emerge when one tries to create a physical geometry from the optimal combination of basis designs. There are two alternative physical interpretations to a superposition of two basis platforms' coefficients. One is that the superposition represents a platform with an *intermediate* shape (e.g. combining two cylinders of different aspect ratios results in a cylinder of intermediate aspect ratio, Figure 4.14). The other is that the superposition represents a platform that *combines* less prominent versions of both platforms (e.g. combining the Spar and Semisub platforms results in a scaled-down spar surrounded by three scaled-down cylinders, Figure 4.15). Intuitively, both of these approaches seem valid for certain ideal cases: for the "intermediate" interpretation, if the two basis geome-

tries are sufficiently similar, and for the "combined" interpretation, if the two basis geometries are sufficiently distant from each other that they do not influence each other's hydrodynamics. Unfortunately, the majority of practical cases fall in between these two extremes.

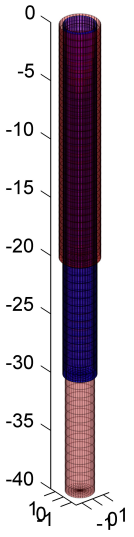

Figure 4.14: Platform geometries showing an "intermediate" physical interpretation (in blue)

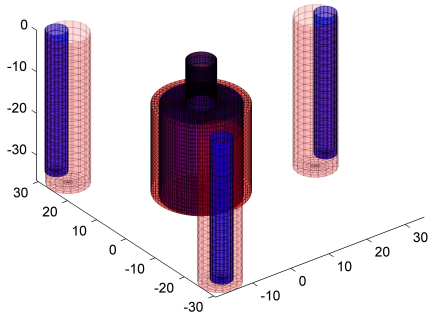

Figure 4.15: Platform geometries showing a "combined" physical interpretation (in blue)

4.5.1 Intermediate Interpretation

The suitability of using linear combination of hydrodynamic coefficients to approximate the performance of a platform geometry that is *intermediate* to two basis geometries can easily be explored for similar geometries. Figures 4.17 and 4.16 show pitch added mass and damping curves, respectively, for three cylindrical platform geometries, each of 150 m^3 volume, but with different aspect ratios. In this exploration, the tallest and shortest cylinders are considered basis designs and the intermediate cylinder is considered a physical interpretation of their combination. The fourth curve in each plot is a 67%-33% superposition of the shortest and tallest cylinders' curves, respectively. This ratio provides the same water plane area as the intermediate cylinder. The resulting coefficient curves are similar to the intermediate cylinder's curves, but are not consistent about which side they err on – whether closer to the shorter

or the taller cylinder. In other words, the parameters change in unique, nonlinear ways as the platform shape is varied. Because of the inherent nonlinearity in how the hydrodynamic coefficients are determined from the geometry, there is no single ratio by which the coefficients of the shorter and taller cylinders can be combined to produce a match of the coefficients of the intermediate cylinder.

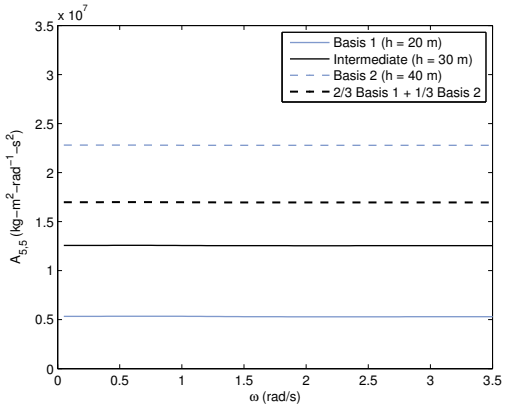

Figure 4.16: Pitch added mass of cylinders

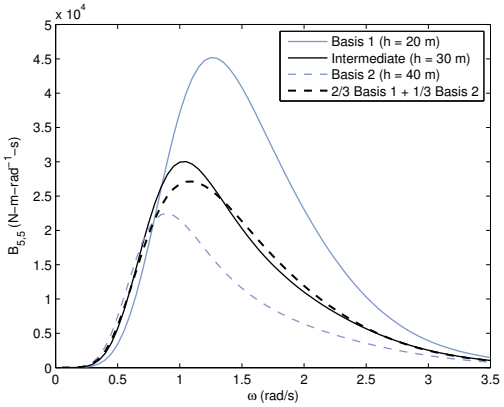

Figure 4.17: Pitch damping of cylinders

4.5.2 Combined Interpretation

A similar attempt for a combination of two very different basis platform geometries can be done. The combination of the Sub and Semisub basis designs is one of the optimal cases for the slack catenary-moored optimization, as was shown in Figure 4.4(b). In this case, it is hard to imagine an intermediate geometry; imagining a *combined* geometry that merges scaled-down versions of both basis designs is more intuitive. Of course, how the two basis geometries are each scaled down is important, and brings up the same nonlinearity problems just described for intermediates between similarly-shaped geometries.

To test the combination interpretation, a platform geometry was designed to try to match the linearly-combined behaviour of 70% Sub and 30% Semisub. The resulting geometry, shown in blue in Figure 4.15, was sized to match the linearly-combined characteristics in terms of water plane area, water plane moment of inertia, displaced volume, and center of buoyancy. This was done by varying the cylinder radii and the spacing of the Semisub's cylinders. For simplicity, the mass of the new geometry was not modeled, but instead assumed to be the same as the linearly-combined mass.

Figures 4.18 to 4.21 compare the hydrodynamic pitch characteristics of the new synthesized geometry with the linearly-combined coefficients it is intended to match. The added mass curve produced by the synthesized geometry falls in between the curves of the two basis geometries, as would be expected, though it is noticeable lower than the curve obtained by the combination of basis geometry curves. The synthesized geometry's damping curve, however, is significantly lower than *both* basis geometries' curves. A similar discrepancy can be observed for the wave excitation curve. The RAO curves in pitch appear to be more reasonable.

The large discrepancies in damping and wave excitation coefficients can be explained in-part by the hydrodynamic interactions between the semisubmersible and submerged vessel platforms when they are combined, which are ignored by simply linearly combining their respective coefficients. For damping, superposition implies *addition* of the damping values. In reality, however, the waves radiated by the Semisub and the Sub during pitch motions will be out of phase with each other at some frequencies, so the damping values should in these cases *subtract* from each other. Therefore, damping curves from a synthesized geometry will be less than those predicted by the basis-function linear combination. For the wave excitation force, it is possible that diffraction effects between the central and outer cylinders act to reduce the wave excitation forces. As the basis-function approach neglects these interactions it would overpredict wave excitation coefficients compared to the more-realistic coefficients calculated from a synthesized geometry. These interaction effects help to explain the results, and demonstrate some important limitations of the basis function approach. Such limitations become more problematic for basis designs having geometries that are close together or intersecting. To overcome these limitations, alternative ways of relating performance and physical geometry will be needed.

4.5.3 Interpretation of Optimization Results

Returning to the optimization results, there is clear evidence that they support the "combined" rather than "intermediate" physical interpretation. Looking at the wave excitation coefficient curves of Figure 4.7 or 4.12, one can see that the magnitude of the optimized (combined) platform's curve is for certain frequencies less than the magnitudes of all the other curves of which it is a linear combination. This is explained by these in fact being complex coefficients, some of them out of phase with each other. This highlights an important driver in the current optimization approach:

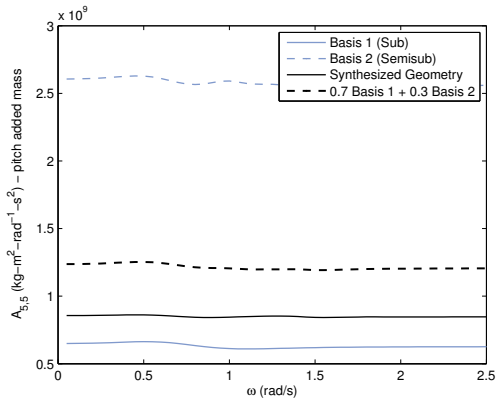

Figure 4.18: Pitch added mass of combined platforms

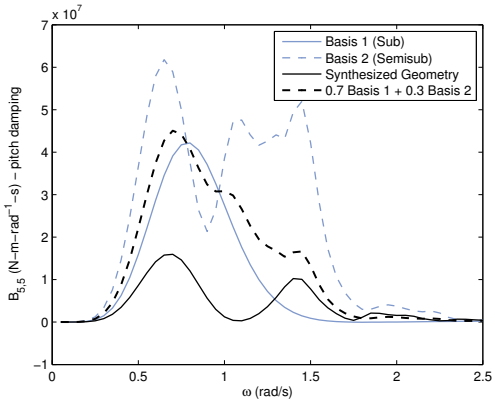

Figure 4.19: Pitch damping of combined platforms

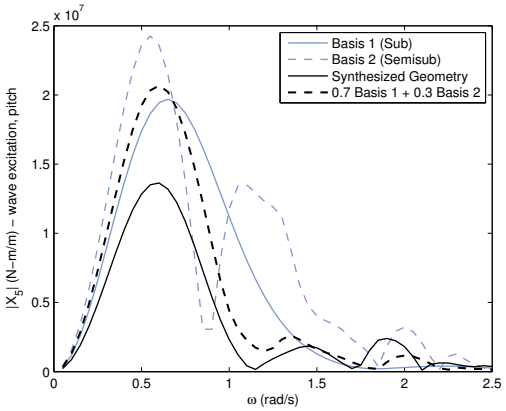

Figure 4.20: Pitch wave excitation of combined platforms

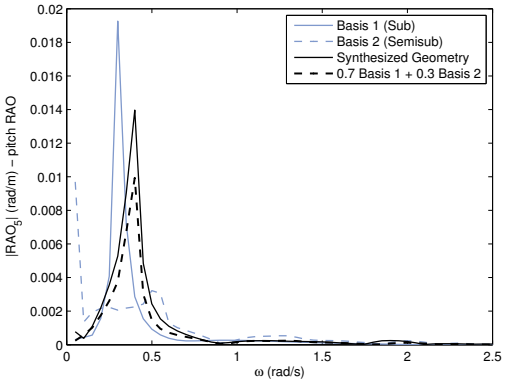

Figure 4.21: Pitch RAO of combined platforms

combining basis platforms with out-of-phase wave excitation can reduce the resulting wave excitation below the values possible for a single platform alone. The excitation reduction occurs because the combined platforms are so different; an intermediate geometry could not achieve the same reductions. It is possible that a platform configuration combining two very different basis geometries (say the Semisub and the Sub) could take advantage of phase cancellation in the excitation forces. However, interactions between the structures would likely emerge in such a combined structure, which cannot be captured in the basis function approach used here.

4.6 Conclusions

The floating wind turbine platform design problem is characterized by a level of complexity that makes methodical and all-inclusive identification of optimal designs next to impossible. This work was a first step in trying to come up with a "hydrodynamics-based" technique that abstracts the geometric details of the platform in order to explore the design space in a simplified, intuitive, and inclusive way. Results from such an approach could then serve as performance guidelines to inform the starting points of more rigorous parametric optimization studies.

A basis function approach was selected to represent the design space as a linear combination of performance characteristics from a set of geometrically-defined "basis" platform designs. A frequency-domain model was used to evaluate the design space, using linear hydrodynamics for the floating platform, and linearized representations of the wind turbine and two alternative mooring systems. This is a low-fidelity modelling approach compared to the time-domain simulations used in Chapter 3. Optimizations were run to find the platform design that minimizes nacelle acceleration for each mooring system, and reasonable results were obtained. The results for a slack catenary mooring system point to submersed volume combined with widely distributed water plane area being design characteristics worth further exploration. The results for the tension leg mooring system support the common TLP-style submersed cylinder as optimal, but with some additional water plane area. In all cases, the semi-submersible configuration is the favoured means of providing distributed water plane area. The semi-submersible (or any multi-cylinder) design has a distinct advantage over most other configurations in that by virtue of it's multiple hulls, the water plane area and moment of inertia can be varied independently. This gives credence to the value of using a semi-submersible configuration with a central cylinder in a more detailed parameter study.

The basis function approach rests on the assumption that linearly combining platform characteristics provides a physical result. The validity of this assumption depends heavily on the specific combination in question. It comes down to physical interpretation. Combinations of coefficients for similarly-shaped geometries can represent an intermediate geometry fairly well. Combinations of coefficients for well-separated geometries can, to a large extent, represent a geometry featuring a combination of each geometry, scaled down. Some simple tests of both interpretations demonstrate their applicability and also their limitations. For combinations of geometries that fall in

between these interpretations, such as intersecting but differently-shaped platforms, interactions between the platforms that cannot be captured by linear combination necessitate the search for a more sophisticated approach.

One option to improve the physicality of the work is to use a more geometry-based basis-function approach that geometrically combines multiple floater styles using basis functions but then realizes the geometric result of the combination and performs the hydrodynamic analysis using WAMIT on that result. This approach would still allow for a wide design space while also ensure that the results have corresponding physical representations. A disadvantage is an increase in processing requirements.

Beyond that, there may be reason to explore other more abstract hydrodynamics-based optimization approaches. These may make use of curve fitting techniques and complex constraints between parameters, or more advanced meta-modelling techniques. Ultimately, striving to come up with models to realistically constrain and relate the performance properties of platforms is an exercise in developing a more formal and comprehensive understanding of the platform design problem.

Chapter 5

Geometry-Based Support Structure Optimization - A Genetic Algorithm-Based Framework

5.1 Introduction

Chapter 4 discussed the merits of a hydrodynamics-based approach to the support structure optimization problem, and explored one of the simplest means of implementing the idea – a basis function approach. Taking that simple approach as far as it would go was an informative exercise, but limitations in the physical interpretability of the formulation were uncovered. There are immense challenges associated with evolving that approach into a robust and physically-valid platform optimization tool or developing an alternative approach within the paradigm of hydrodynamics-based optimization. While doing so would be an interesting academic exercise, in the interest of practicality, attention was turned to making contributions in the more proven realm of geometry-based optimization approaches.

To provide a global optimization framework for the support structure, a parameterization that captures the full range of the design space is required. This means that the parameterization must be able to represent existing design geometries as well as feasible not-yet-conceived ones. Creating a scheme that meets those requirements is one challenge. A second challenge is integrating models of the wind turbine, float-

ing platform, and mooring line dynamics into a combined model of the full floating wind turbine system in order to evaluate the designs created by the scheme. A third challenge is creating the optimization framework that the scheme operates in; the many discontinuities in the design space (from different numbers of hulls, different configurations, etc.) and potential for multiple competitive local optima require a special type of optimization algorithm. As well, careful arrangement of the design variables is important in order to provide a well-organized design space.

An optimization framework that meets these challenges has been developed. It features a more flexible geometry parameterization and a more sophisticated design space exploration approach in order to span a greater extent of the design space than existing support structure optimizations in the floating wind turbine literature [16, 56]. The computation time for an optimization is in the order of days. The framework has three components:

1. a support structure design variable scheme that provides the parameterization to describe the design space (discussed in Section 5.2),

2. a frequency-domain dynamics model to evaluate points in the design space (discussed in Section 5.3), and

3. a genetic algorithm to manage the exploration of the design space (discussed in Section 5.4).

Each of these components was developed in sufficient depth in order to be able to demonstrate the operation of the overall framework, as is done in Section 5.5. As is discussed in Section 5.7, there is room for improvement in each component, such as removing restrictive assumptions from the parameterization, improving approximations in the dynamics model, and speeding up the optimizer. Focussing on these refinements was considered unfruitful at present given the need for hard-to-obtain site data and cost data in order to produce accurate design optimization results. Lacking these data sets, input values and cost functions are designed first and foremost to demonstrate the *potential* of the framework. The primary contribution is the integration of the three specialized components – parameterization scheme, dynamics model, and genetic algorithm – into a unified and adaptable optimization framework.

The goal of this tool is not to automate the design process – the design problem is too complex and multifaceted for that. Rather, the intention is to provide a framework that can be applied to a given siting scenario to produce a list of the most promising

floating support structure configurations. These configurations can then serve as starting points for more detailed design processes. This way, more conventional design approaches (and optimizations) will converge to optimal designs faster, and promising design options will not be overlooked for lack of imagination. As well, application of the framework may provide insight into the nature of the design space.

5.2 Support Structure Parameterization

The heart of the optimization framework is the support structure parameterization scheme. In the interests of optimization simplicity, the scheme was made with the aim of describing the widest range of feasible platform and mooring system configurations with as few design variables as possible. Additionally, preserving some degree of continuity or consistency in the effects of the design variables across their ranges and avoiding any redundancy (in which similar configurations could be achieved by different combinations of design variables) is important in order to keep the design space organized. Adhering to these objectives speeds up optimization convergence and makes results easier to interpret.

The scheme features nine design variables, and consists of components to deal with:

- the platform geometry,
- the mooring line configuration,
- the size of tendon arms for taut mooring systems,
- the size of structural elements connecting the platform cylinders,
- the ballast and mass properties of the platform, and
- the cost of the overall structure.

5.2.1 Platform Geometry

A scheme based on vertical cylinders was selected to parameterize the platform geometry. This was deemed the most straightforward option considering the range of existing platform designs and the preference for cylindrical hulls for hydrodynamic, structural, and manufacturing reasons. The scheme consists of a central cylinder

whose radius and draft are variable, with the additional control of a variable amount of taper near the water plane, as well as an array of three or more outer cylinders whose radius, draft, and distance from the center are collectively variable. Because heave plates are common for multi-cylinder platforms, the outer cylinders can feature circular heave plates of variable size at their bases. The inner cylinder or the outer cylinders can be excluded by using a minimum threshold in their radius design variables. This is an example of implementing a discontinuous feature in a continuous design variable in an intuitive way, though a gradient-based optimizer would still be challenged by this scheme. The tapered section of the central cylinder is set to occur from 1/4 draft to 1/8 draft. Figure 5.1 illustrates the geometry scheme, before mooring lines, connective structural elements, and structure above the water plane are added. The eight parameters of the geometry scheme are provided in Table 5.1.

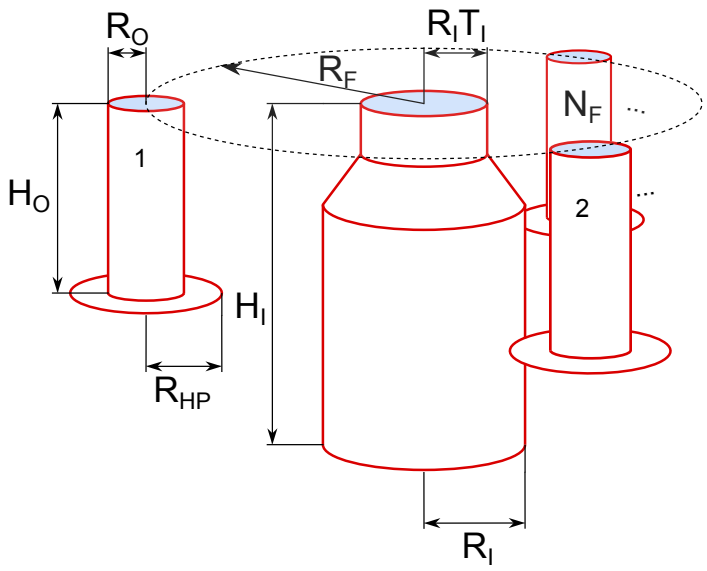

Figure 5.1: Vertical cylinder-based platform geometry scheme

Some constraints are applied to these variables to reflect design requirements or pragmatic modelling limitations:

- The top diameter must not be less than the tower base diameter (6 m [24]) to provide sufficient strength for supporting the tower. If there is no central cylinder, this constraint is applied on the outer cylinders, to facilitate the turbine being mounted on one of them.

- A minimum of 1 m of separation must exist between all cylinders in the design,

Table 5.1: Platform geometry scheme design variables

Description	Variable	Min.	Max.
inner cylinder draft	H_I	2 m	150 m
inner cylinder radius	R_I	3 m	25 m
inner cylinder top taper ratio	T_I	0.2	2
number of outer cylinders	N_F	3	8
radius of outer cylinder array	R_F	5 m	40 m
outer cylinders draft	H_O	3 m	50 m
outer cylinders radii	R_O	1.5 m	10 m
outer cylinders heave plate radii	R_{HP}	0 m	20 m

including accounting for heave plates, to avoid numerical modelling problems and maintain physical plausibility.

- The permitted amount of taper of the inner cylinder radius is restricted for drafts less than 10 m. The variation is linear, reducing to no taper ($T_I = 0$) at the minimum draft ($H_I = 2\,\text{m}$). This is to prevent the changes in water-plane area and bottom-slamming that could happen if waves amplitudes reach the taper draft.

A number of conditions are built into the geometry scheme to represent the various discontinuities that may arise:

- Inner or outer cylinder radii that are below the respective minimum bounds ($R_I < 3\,\text{m}$ or $R_O < 1.5\,\text{m}$) signal that the respective cylinder(s) do not exist in the platform. In that case, the remaining values associated with the cylinder (eg. draft) are assigned a NaN value to be properly handled by the optimization algorithm.

- If the heave plate radius is less than the outer cylinder radius, heave plates are not considered in the analyzed geometry.

Using these conditions allows the generic geometry parameterization shown in Figure 5.1 to produce a wide range of platform configurations (see for example Figure 5.18(a).

5.2.2 Mooring System

The mooring system scheme adds one additional design variable, x_M, to the design space, for a total of nine design variables. The mooring line configuration in the

framework is then determined by this variable in conjunction with several of the platform geometry design variables and the water depth. Mooring systems for floating wind turbines conventionally fall into one of two distinct configuration types – slack catenary lines or taut vertical lines – with variations within these configurations occuring in the number of lines, the fairlead and anchor locations, and the line lengths (or, equivalently, tensions) [1]. Intermediate "taut-catenary" configurations also exist, as have been investigated by researchers at MIT [51].

As illustrated in Figure 5.2, the non-dimensional mooring configuration design variable, x_M, controls a mooring configuration algorithm that transitions smoothly between a taut vertical TLP configuration and a slack catenary configuration with widely-spaced anchors. The algorithm calculates the number of mooring lines and the fairlead locations based on the platform geometry and x_M. This relatively constrained arrangement was used to avoid wasting computation time on impractical mooring systems.

If there is only one central cylinder, three (if slack) or four (if taut) lines are used and they connect at the cylinder bottom circumference for a taut mooring system and half way along the draft for a slack mooring system. If there are multiple cylinders, the mooring system depends on the comparative drafts of the inner and outer cylinders. For a central cylinder whose draft is more than double the draft of the outer cylinders, the mooring system is attached to the central cylinder, identical to the case of there being no outer cylinders. Otherwise, the number of lines is equal to the number of outer cylinders and the fairleads are located at the outer edge of the bottom of each of the outer cylinders.

The anchor locations are determined by x_M and vary linearly with x_M from lying directly under the fairleads (when $x_M \leq 0$) to having a horizontal spread of double the water depth (at $x_M = 2$). The algorithm that determines the length and tension of the mooring lines is set up such that the transition between taut and slack configurations occurs when the anchors are spaced a horizontal distance from the platform equal to the water depth ($x_M = 1$).

For slack mooring configurations, the unstretched mooring line length is deter-

[1] In this work, "slack" implies that portions of the mooring lines rest on the sea floor, hence the only vertical force imparted to the platform by the mooring lines is from the lines' weight. "Taut" implies that no part of the mooring lines lies on the sea floor, meaning the vertical force on the platform comes from both the weight of the mooring lines and their pretension.

mined according to

$$L_{unstr.} = \sqrt{l_x^2 + l_z^2} + \frac{l_z}{12} \tag{5.1}$$

where l_x is the horizontal distance from anchor to fairlead and l_z is the vertical distance from anchor to fairlead. The last term in this equation was chosen to give line lengths similar to those of the OC3 Hywind design.

For taut configurations, the mooring line length is chosen such that the resulting line tension cancels any surplus buoyancy in the system (i.e. it is a function of the platform decision variables) and no ballast in the platform is assumed. Taut vertical mooring configurations where the lines are held at a distance from the platform cylinder(s) by horizontal "tendons" (see Figure 5.10(c)) are supported by negative values of x_M. The length of these tendons is then equal to $(-50\,\mathrm{m})x_M$.

At present, the mooring line material properties are kept fixed, with elasticity modulus of 6 MPa and density of 12200 kg/m^3. The mooring line cross-sectional area is varied inversely to the number of lines to keep the mooring system total mass proportional to the individual line length only. For three lines, the diameter is 90 mm, consistent with the OC3 Hywind design [55].

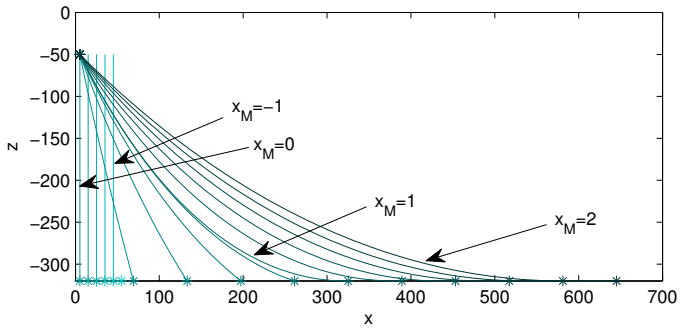

Figure 5.2: Demonstration of mooring line layouts generated by mooring algorithm for x_M values varying from -1 to 2

5.2.3 Taut-Mooring Tendon Arms

If a taut mooring configuration is used with tendons holding the fairleads at a distance radially from the platform, a scheme is needed to assign realistic properties to these tendons. The tendons can have important effects on the mass, buoyancy, hydrodynamics, and cost of the support structure. To model the necessary properties, these horizontal members are treated as steel tubes with a constant wall thickness

to radius ratio of $k = 5\%$. Their diameter is chosen for a bending moment criterion, based on the bending moment at the cylinder connection point. This provides a strength-based constraint on the length of the tendons, L. The load considered is the vertical component of the maximum steady-state mooring line tension multiplied by a safety factor of $F_S = 3$. This safety factor was chosen to calibrate the scheme to the tendon diameter of the MIT/NREL TLP. A yield stress for steel of $\sigma_y = 200\,\mathrm{MPa}$ is used.

The bending stress at the root of the tendon is calculated as

$$\sigma = \frac{My}{I} = \frac{8F_S F_{line,z} L}{\pi k D^3}, \tag{5.2}$$

where $F_{line,z}$ is the vertical component of the steady mooring line tension on the windward fairlead at the maximum wind thrust load, corresponding to the maximum steady-state tension at the fairlead. The mass of the tendon elements, $m_{tendons}$ therefore has the following proportionalities:

$$m_{tendons} \propto N_{lines} L D^2 k \propto N_{lines} L^{(5/3)} F_{line,z}^{(2/3)} k^{(-2/3)}. \tag{5.3}$$

With the tendons sized, their mass, buoyancy, and hydrodynamic contributions are then determined.

5.2.4 Float-Connecting Truss Members

If the platform features multiple cylinders, the structure connecting them together is an important contributor to the support structure dynamics and cost. Indeed, the high cost of such structures is what deters against platforms large enough to support multiple wind turbines [2]. In this framework, the structure connecting multiple hulls is modelled as a truss segment consisting of three tubular members – two horizontal beams and one diagonal crosspiece – between each pair of connected cylinders, as shown in Figure 5.3. For platforms without an inner cylinder, adjacent cylinders are connected; for platforms with both inner and outer cylinders, each outer cylinder is connected to the inner cylinder. For strength, the truss section is kept quite tall, with the bottom member at 90% of the inner or outer cylinder draft (whichever is less) and the top member at a height of half the airgap above the waterline.

The three members are treated as hollow cylinders with a fixed wall thickness to

[2]For an example of multi-turbine support structures, see [57].

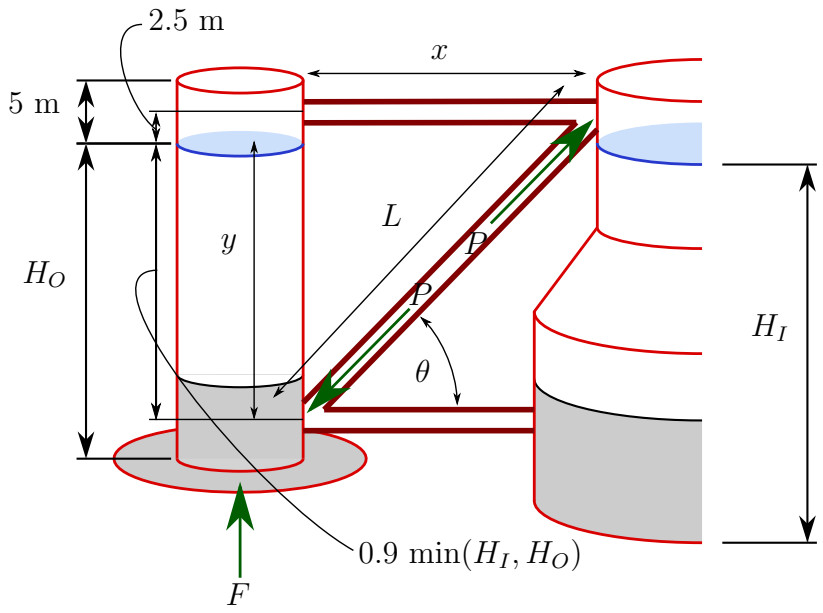

Figure 5.3: Truss scheme for connecting cylinders

radius ratio of 5%. The diameter of all three is chosen based on the pinned-pinned critical buckling load, P_{crit}, of the diagonal member:

$$P_{crit} = \frac{\pi^2 EI}{L^2}, \tag{5.4}$$

where L is the length of the member, $E = 200\,\text{GPa}$ is the elasticity modulus of steel, and I is the tubular section's moment of inertia. The compressive load, P, on this member is calculated based on a vertical load on the truss equal to the displaced weight of one of the outer cylinders, $\rho V g$, or the maximum steady-state mooring tension, $T_{line\,max}$, if the mooring lines are connected to the outer cylinder, whichever is larger:

$$P = F_S \frac{\max(\rho V g, T_{line\,max})}{\sin(\theta)}, \tag{5.5}$$

where θ is the angle of the diagonal member. A safety factor of $F_S = 10$ is applied to this vertical load; this value was found to produce truss member diameters that are similar to those used for the OC4 WindFloat design.

The diameter of the members can be calculated as

$$D = \left(\frac{8}{\pi^3}\frac{PL^2}{Ek}\right)^{1/4} = \left(\frac{8}{\pi^3}\frac{P(x^2+y^2)^{3/2}}{EKy}\right)^{1/4}. \tag{5.6}$$

The mass of the truss, m_{truss}, is then proportional to the following:

$$m_{truss} \propto N_F LA \propto N_F LD^2 k \propto N_F L^2 P^{(1/2)} k^{(1/2)}, \tag{5.7}$$

where A is the cross-sectional area of the member.

The justification for this approach is that in operating conditions, particularly if active ballast is used, it is entirely possible that the vertical force from an outer column on the system will reverse, imposing a compressive load on the horizontal members; sizing the horizontal members based on the buckling load sizing of the diagonal members is therefore an approximate way of designing for this potential load reversal. Alternatively, if small outer cylinders are used with a large inner cylinder, the forces of the mooring lines attached to the outer cylinders will make up the majority of the load on the trusses, so they should be sized for this load accordingly - using the maximum steady-state tension is simply a way to approximate the correct proportionality.

This scheme enables treatment of the changes in platform mass, buoyancy, hydrodynamic properties, and cost as a result of the connective elements.

5.2.5 Platform Mass and Ballast

The platform geometry scheme and mooring system scheme go hand-in-hand with a mass model that predicts the mass characteristics of the platform and determines the use of ballast. The main component of the distributed mass of the platform is the mass of the cylinders, modelled by assuming constant-thickness steel on their surface areas, including above the waterline. This thickness is greater than that of physical designs to also represent the mass of structural elements (bulkheads, stiffeners, stringers, etc.) within the platform. The other contributions to the platform mass come from the heave plates, the connective trusses and taut-mooring tendons, and the ballast. The ballast and connective structure are shown in Figure 5.4.

In the results here, the heave plate steel thickness is taken to be 30 mm and the hull steel thickness is determined by $T_{hull} = 50\,\text{mm} + 0.0003\max(H_I, H_O)$ to account for the increased pressure forces at greater water depths. These values were selected in order to calibrate the scheme's mass model to be able to recreate the mass properties

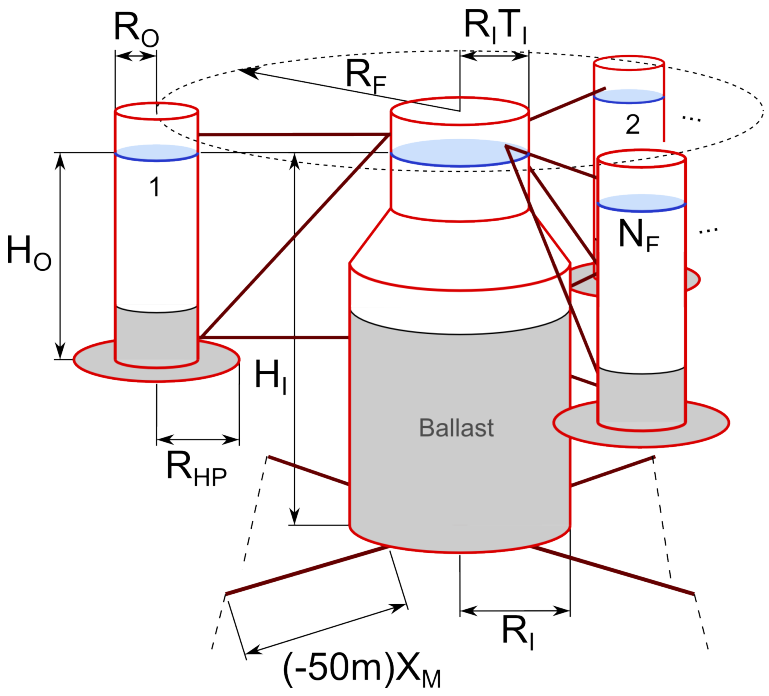

Figure 5.4: Platform geometry scheme with ballast and connective structure

of the Hywind and WindFloat designs.

Once the mass of all the structural components is known, the amount of ballast is determined. The ballast mass is set according to the surplus buoyancy of the system - the remaining buoyancy force after subtracting wind turbine weight, platform structural weight, and the vertical component of mooring line tensions. In the case of taut mooring line configurations ($x_M < 1$), no ballast is applied and instead any surplus buoyancy is taken up by increasing the mooring system tension. In the case of slack mooring systems, ballast is added from the bottom of the deepest cylinder(s) upward, to a common top level across all cylinders. The contribution of the ballast to the distributed mass of the platform can then be calculated from these volumes.

The possibility of active ballast – pumping water ballast between cylinders to counter steady overturning moments – is also included for multi-cylinder designs. The effect is only considered in the analysis of the platform's static pitch angle; the shift in the platform center of mass is neglected in the dynamic analysis. For the static pitch angle calculation (which is described in Section 5.4.3), the moment available from ballast shifting is modelled as

$$M = R_F \frac{N_F m_{ballast\,O}}{2} g \tag{5.8}$$

where $m_{ballast\,O}$ is the mass of ballast assigned to each outer cylinder (before shifting). This equation corresponds to shifting the ballast proportionally to the x-axis location of each cylinder (as shown in Figure 5.5) and is a relation that holds true independent of the number of floats.

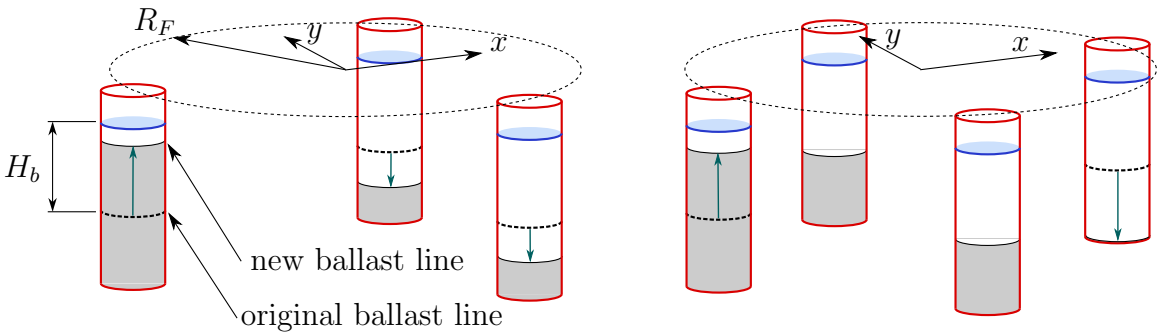

Figure 5.5: Ballast shifting strategy

The choice of ballast material is important. On one hand, a dense non-adjustable material such as gravel or concrete provides the high densities needed for economical ballast-stabilized designs. On the other hand, using water as ballast allows the active ballast-shifting that helps wider buoyancy-stabilized platforms stay level in strong winds. For the results here, concrete ballast with a density of 2400 kg/m^3 is used, in order to enable ballast-stabilized designs. The ramifications for buoyancy-stabilized designs are of lesser importance, because the shallow draft of these designs results in a relatively small change in center of gravity compared to the situation where water ballast is used[3].

A certain amount of structure is needed above the waterline. Common freeboard or airgap heights range from 5 m in Tracy's parameter study [51] to 10 m in the WindFloat design [25]. A height of 5 m is used in the results generated here.

5.2.6 Support Structure Costs

The stability of a floating structure generally improves with structure size. The cost of the structure is the main factor that constrains this. As such, accounting for the support structure cost is crucial for a realistic representation of the design problem. The installed cost of the system is modelled as a combination of three component costs – for the floating platform, the mooring lines, and the anchors.

[3]The contention is that if water ballast was instead used in the framework, the effect on buoyancy stabilized designs would not significantly alter the results, with a small enough change in the center of mass to have only a minor effect on the platform's metacentric height.

Platform Structure

The cost of the platform structure is considered proportional to the structure mass. This accounts for material costs as well as fabrication and installation costs in the simplest way. Specific costs pertaining to different structural components are neglected. In a floating wind turbine platform study published with different cost numbers for different components, the per-mass cost differences between columns, trusses, braces, and deck differ by no more than 20% of the mean value across these different components [58]. Considering the cost numbers of that study as well as the per-mass material and fabrication costs presented in [38], a cost of $2.50 per kg of platform is used in the results here. Because of the inexpensive materials that can be used as ballast, a ballast cost is not used.

Mooring Lines

The cost of the mooring lines is treated as a linear function of their total combined length and the maximum steady-state tension they have to withstand. This approach defines the cost proportionally to the mass of the lines, because the cross-sectional area required is proportional to the line tensile strength. Implicit in the approach is the assumption that the peak tension the lines need to withstand will be proportional to the steady-state tension at the maximum wind speed condition. This is one of many approximations made to simplify the evaluation procedure of the framework. Nonlinearities in the cost function, which could arise from the inclusion of installation costs, line purchase costs, etc. are neglected for simplicity.

The line cost is based on a factor of $0.42 /m-kN which is multiplied by the total line length and the maximum steady-state line tension. This gives final line cost results that fall within the range of costs spanned by [38], [51], and [58].

Anchors

Anchor cost, for which installation cost is a significant component, is affected by both discrete anchor technology options and continuous anchor size factors. A three-technology anchor cost model was used in the framework to provide a simplified treatment of the anchor cost factors.

The choice of anchoring technology is determined primarily by the direction of loading on the anchor. The three most applicable anchor types are drag-embedment anchors, vertical-load drag-embedment anchors (VLAs), and suction piles:

- Drag-embedment anchors work by penetrating in the seabed as they are dragged by a horizontal force from the mooring line. They can penetrate to great depths in this way and have holding capacities of over 50 times the anchor weight depending on soil conditions [59]. Large loads on the anchor will further embed the anchor rather than pulling it out, provided the mooring line angle is shallower than 10-20 degrees (depending on soil conditions) [60].

- Vertical load anchors (VLAs) are similar to drag-embedment anchors and are installed in the same way. However, after being penetrated to the required depth, a shear pin causes the angle of the fluke to increase so that the mooring line tension is nearly normal to the fluke. This maximizes the anchors holding capacity for normal loads. VLAs suffer a weakness that movement resulting from sustained vertical load can result in reduced holding strength and eventually cause the anchor to come loose [61]. This necessitates a higher safety factor for VLAs and a more thorough site selection and installation process, adding cost [62]. Bruce anchors therefore recommends their current near-normal load anchor, the Dennla, for line angles only up to 45 deg [63].

- Suction piles are open-bottom piles that are embedded into the seafloor by pumping water out of them through a sealed top. The resulting suction drives the anchors into the seabed and provides a strong resistance to vertical loads, making them well-suited to taut-moored configurations.

The discrete anchor options, with their different installation procedures and load angle tolerances, necessitate a discontinuous cost function. Accordingly, different coefficients are used for each technology. In the framework, the angle of the mooring lines at the anchor (which is a function of x_M) determines which anchor technology will be used. The cost of the anchors is modelled as a linear function of the maximum steady-state load on the anchors. As with the mooring lines, the anchors are sized based on steady rather than peak loads because the latter would require a more involved iterative design approach. A fixed per-anchor installation cost is also included. The anchor costs and line angle criteria employed in the framework are given in table 5.2.

Table 5.2: Anchor cost model

Anchor Technology	Line Angle	$/anchor/kN (line tension)	$/anchor (installation)
drag embedment	0°-10°	100	5000
vertical load (VLA)	10°-45°	120	8000
suction pile	45°-90°	150	11000

5.3 Modelling and Evaluation Methodology

After the support structure decision scheme produces a design, the performance of the design needs to be evaluated. The evaluation of each point in the design space and calculation of its objective function value are handled by a six-DOF frequency-domain model created in Matlab, an expanded version of the model used in Chapter 4. This linear model provides a computationally-efficient way of coupling the dynamics of the wind turbine, mooring system, and floating platform. Loads from steady winds and regular (monochromatic) waves are included. The DOFs considered are the six rigid-body modes of the platform. By definition, the frequency-domain model assumes that the platform motions are at the same frequency as the incident waves and that the incident waves are regular. While this means that the transient response of the system cannot be modelled, the assumption of linearity implies that the responses at different wave frequencies can be superimposed according to a wave spectrum to predict the system behaviour in irregular sea states, as was done in equation (4.6).

The combination of models from which the linear coefficients are generated shares many commonalities with the combination used in the basis function work described in Section 4.3. The major differences relate to the three new design features in the GA-based framework – heave plates, truss and tendon members, and a parameterized mooring system – and the fact that all coefficients except for those of the wind turbine now have to be calculated anew for every design point.

Figure 5.6 lists the loads acting on the floating platform that are considered in the frequency-domain analysis, with loads calculated from external models indicated with parentheses.

5.3.1 Platform Hydrodynamics

In a departure from the model of Chapter 4, the inviscid linear platform hydrodynamics calculated by WAMIT are supplemented with linearized calculation of the viscous

- Wind turbine loads:
  - linearized mass, damping, stiffness (FAST)
  - thrust (FAST)
  - weight
- Platform static loads:
  - mass
  - weight
  - buoyancy
  - hydrostatic stiffness (WAMIT)
- Platform hydrodynamic loads:
  - added mass (WAMIT)
  - wave-radiation damping (WAMIT)
  - linearized viscous drag
- Mooring system loads:
  - static force offsets (quasi-static model)
  - linearized stiffness (quasi-static model)

Figure 5.6: Loads on floating platform

drag forces on the platform cylinders(s), truss and tendon members, and heave plates. The linearization of these viscous drag terms is done iteratively during the equation-of-motion solution because the linearized terms are amplitude-dependent.

Linear Hydrodynamics

To calculate the linear hydrodynamic loads on the platform, the required linear hydrodynamic coefficients are generated for each platform design by the panel method code WAMIT. Before WAMIT is called, a meshing routine created in C++ discretizes the surface of each candidate platform design, including the heave plates, and generates the WAMIT geometry file. The same C++ routine also performs the platform mass calculations and handles the calls to the mooring line model. This process is embodied in a DLL and interfaced to Matlab using a .mex file. The interface returns the aggregate platform properties to the Matlab-based frequency-domain model of

the combined system.

Heave Plate Viscous Drag

The viscous drag of the heave plates, which is quadratic in nature and not related to wave radiation damping, is not modelled by WAMIT's linear potential flow method so another model is required to capture it. The most common approach is to use the drag term of Morison's equation (2.2) where the viscous drag force, f, is modelled as the product of a drag coefficient, C_D, the heave plate area, A_{hp}, and the square of the normal component of the relative fluid velocity, $u(t)$.

$$f(t) = \frac{1}{2}\rho A_{hp} C_D u(t)|u(t)|. \tag{5.9}$$

To fit into a frequency-domain model, this relation needs to be linearized. Tao and Dray provide such a linearization [64]:

$$f(t) = \frac{2}{3}\rho D^3 \omega B' u(t), \tag{5.10}$$

with B' being a function of KC (the Keulegan-Carpenter number) for which empirical relations have been developed. The relation used here is $B' = 0.2 + 0.5K_C$, as taken from [64].

The approach of equation (5.10) was implemented into the combined frequency-domain system model. Wave kinematics are not included in the calculation of relative fluid velocity for simplicity, on the grounds that heave plates are at a depth where wave velocities are quite low. The model provides the viscous drag contribution to the platform damping in the three DOFs most affected by the heave plates - heave, pitch, and roll. Viscous drag contributions in other DOFs are ignored because these DOFs cause lateral motion of the heave plates for which the damping forces are smaller and harder to model. The wave-radiation damping and added mass from the heave plates are provided by the WAMIT analysis.

Platform Viscous Drag

A viscous drag model was also implemented for the platform cylinders, the submerged connective trusses, and the taut-mooring tendons, since the damping forces on slender cylinders are not adequately accounted for by a linear hydrodynamics approach alone. For these elements, a linearization of the drag term of Morison's equation (2.2) is used,

with a drag coefficient of 0.6. The linearization is taken from the work of Borgman [65] as described by Savenije and Peeringa [66]:

$$df_{drag,lin.} = \left(\sqrt{\frac{8}{\pi}} \sigma_u \frac{1}{2} \rho D C_D u \right) dL. \tag{5.11}$$

where $df_{drag,lin.}$ is the drag force on a cylindrical section of length dL and diameter D, and σ_u denotes the standard deviation or root-mean-square (RMS) of u. It should be remembered that these supplementary linearized drag calculations are only for transverse forces on the cylindrical elements – u is the transverse component of the relative water velocity.

By accounting for viscous drag forces from lateral motions of the platform cylinders in a strip-theory approach (where the cylinders are discretized axially into strips) a viscous damping contribution can be calculated for surge, sway, roll, pitch, and yaw DOFs. Because the connective trusses and tendons are not in parallel, the strip theory approach applied to them yields viscous damping contributions in all DOFs.

As was done with the heave plate damping, wave kinematics are neglected and only the structure motions are used in the calculation of velocity; this is necessary to maintain the linear frequency-domain representation of the problem.

Connective Element Added Mass and Drag

To avoid the complexity of having to create a panel mesh for the connective trusses and tendons, and because the slenderness of these components makes their wave-radiation contributions relatively small, the trusses and tendons are not included in the WAMIT analysis. Rather, their hydrodynamic properties are accounted for by the viscous drag linearization already mentioned, and an added mass calculation of the form used in Morison's equation. The added mass coefficient used is 0.97 and the damping coefficient is 0.6 – these are the same coefficients as were used in the OC3 Phase IV modelling and in the modelling of the original Hywind design [67]. Only the bottom members of the trusses are modelled. The top members are above the water line. The diagonal members are difficult to model because their wetted area is variable and they would be subject to significant wave velocities as they pass through the water plane. Furthermore, there are many options for the layout of these diagonal members and the approach taken here is fairly arbitrary. In light of these factors, the hydrodynamics of these elements are neglected and their value to the framework is

confined to accounting for contributions in buoyancy, mass, and hydrostatic stability.

The linearization of the drag on each discrete cylinder section (as in equation (5.11)) is done with respect to the transverse component of its RMS velocity (σ_u) calculated across the full frequency spectrum using the approach described in Section 2.3.4 [4]. Because the linearization depends on the amplitude of motion, it is interdependent with the solution of the equation of motion for the system. Therefore, an iterative approach is taken in solving for the motions of the platform, with the viscous drag linearization for each cylinder component updated at each iteration.

5.3.2 Wind Turbine

As was done in Chapter 4, a linear representation of the wind turbine is used, with linearized coefficients obtained using FAST's linearization functionality for each wind speed condition. To limit the complexity of the model, these linearizations are generated at a fixed static pitch angle rather than one that is adjusted for each platform at the thrust load of each wind speed. A value of zero pitch was chosen because many platforms pitch very little or use techniques such as active ballast to eliminate significant static pitch angles. The thrust load at each wind speed determines the lateral load that the mooring system needs to resist and hence it also determines the static surge offset of the system in equilibrium and the behaviour of the mooring system at that position.

5.3.3 Mooring Lines

While Chapter 3 revealed some important accuracy limitations of quasi-static mooring line models, these inaccuracies are not out of proportion to the level of approximation inherent in the frequency-domain modelling approach. Furthermore, representing the dynamic effects of the mooring lines in the frequency domain would be very difficult and not within the computational budget of the optimization framework. The generation of mooring line stiffness matrices is therefore handled by a quasi-static mooring line model.

[4] An alternative approach is to linearize with respect to the velocity at each frequency segment, calculating a different damping coefficient for each frequency. Because the aim is to model the overall response in irregular wave conditions, the approach of using the RMS velocity from combining the responses across the frequency spectrum is thought to be more appropriate. It is hoped that some results obtained from some recent scale model wave energy converter testing will lend experimental support to this.

The model used in this case is Catenary, the mooring line subroutine of FAST, which has been translated into C++ to provide easy integration with other parts of the framework that are also implemented in that language. The linearization routine accepts the decision variables, water depth, and wind speeds, applies the mooring configuration algorithm discussed in Section 5.2.2 to determine the mooring configuration, and uses several layers of iterations and perturbations when calling the quasi-static mooring line model to obtain linearized mooring stiffness matrices. A separate matrix is generated for the displaced equilibrium position of the platform corresponding to the wind turbine thrust load at each wind speed being analyzed. As noted by Tracy [51], the effect of the mooring system changes as the thrust force from the wind causes the platform to be displaced from its centered position. Therefore, for the level of detail desired in the framework, it is important to linearize the mooring system reactions about these displaced equilibrium positions.

5.4 Genetic Algorithm Optimizer

A GA optimizer was selected as the most flexible and straightforward way to systematically explore the design space, given the potential for multiple local optima and the diverse, interrelated, and often-discontinuous design variables. A pure gradient-based approach across the full configurational design space would have been defeated by the discontinuities present. While more sophisticated hybrid metaheuristic optimization approaches that incorporate gradient-based methods around local optima could be more efficient, the simpler GA approach was retained to avoid excessive development effort and complexity in the optimization algorithm. The GA also provides a family of locally-optimal designs, which is more useful than a single optimal design at this level of model fidelity for gaining insight into the characteristics of the design space.

5.4.1 Cumulative Multi-Niching Genetic Algorithm

The algorithm developed specifically for this framework, referred to as the Cumulative Multi-Niching (CMN) GA, has two goals that are not common to all GAs: to use as few objective function evaluations as possible and to be able to identify and converge to multiple local optima. Because the evaluation of each individual design (and the hydrodynamic analysis in particular) is vastly more time consuming than the operations of the GA itself, the algorithm has features designed to limit redundant or

unproductive fitness function evaluations. With the possibility of multiple local optima in the design space, the algorithm is also designed to support multi-niching - the ability to converge to multiple optima simultaneously. Recognizing that the limited fidelity of a frequency-domain model may create distortions in the design space and that additional factors not included in the framework also affect the choice of optimal design, the algorithm was developed to explore local optima in an equitable way, regardless of the comparative fitness values of the local optima, so that potentially promising configurations are not discriminated against based on potential deficiencies in the submodels. Choosing a final optimized design would then require two steps beyond what is provided by the framework:

1. Confirm the performance of the local optima using higher-fidelity time-domain simulation tools, which are necessary in any case for verifying, fine-tuning, and elaborating on the optimal-design results.

2. Apply additional design considerations – such as installation requirements, serviceability, and aesthetics – to complete the decision matrix.

The CMN GA has a unique arrangement of features. The most distinctive feature is that it is cumulative; each successive generation adds to the overall population. By never discarding individuals from the population, the GA can make use of the information from every objective function evaluation as it explores the design space. (Even unfit individuals are valuable in telling the algorithm where not to go.) The key to making the cumulative approach work is the use of an adaptive proximity constraint that prevents offspring that are overly similar to existing individuals from being considered. By using a distance threshold that is inversely proportional to the fitness of nearby individuals, this proximity constraint encourages convergence around promising regions of the design space and allows only a sparse population density in less-fit regions of the design space. An example of this effect is shown for a contrived two-dimensional design space in Figure 5.7.

These distinctions from other GAs enables a number of unique features in the genetic operations of the algorithm that together combine to make the cumulative multi-niching approach work. The selection and crossover operations use a novel technique that considers both proximity and fitness in choosing mates for crossover, providing the required multi-niching capability. A fitness scaling operation makes the GA treat local optima equally despite potential differences in fitness. The algorithms

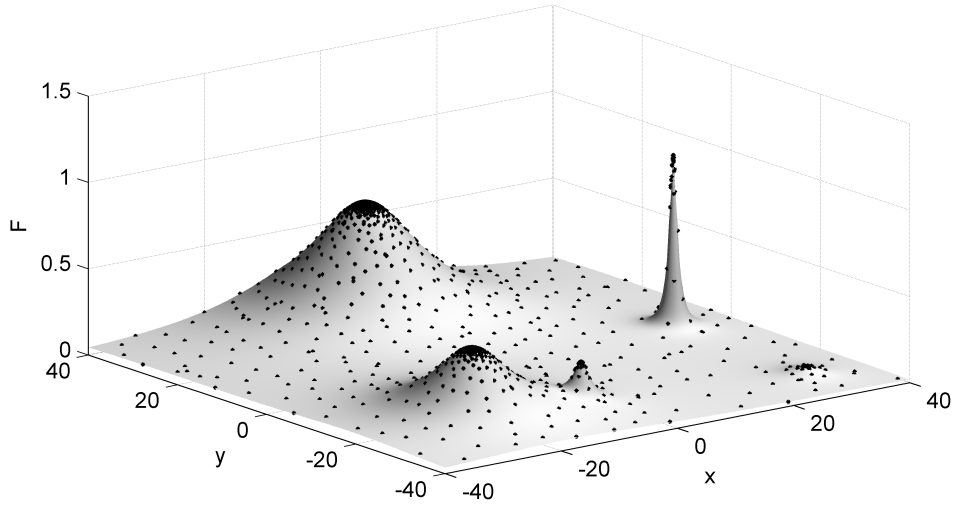

Figure 5.7: Design space exploration of the CMN GA on a sample two-variable objective function (function F4 in [68])

arrangement of genetic operations provides fast and robust convergence to multiple local optima. Benchmark tests alongside three other multi-niching genetic or evolutionary algorithms show that the CMN GA has a greater convergence ability and provides an order-of-magnitude reduction in the number of objective function evaluations required to achieve a given level of convergence. These results, and a more thorough description of the genetic algorithm are provided in the paper in Appendix A.

For the present results, the GA was terminated once a target population size was reached – either 1000 or 1500 depending on the situation. These numbers were chosen to provide sufficient opportunity for the population to converge around the significant local optima. For a given optimization problem, the GA population will tend to converge after a consistent population size has been reached, beyond which only minuscule improvements in fitness will be achieved.

5.4.2 Optimization Objectives

While minimizing cost of energy (COE) is the overall optimization goal for a renewable energy technology, a simpler less-general optimization problem formulation can be used for the floating wind turbine support structure design problem, in order to avoid the additional considerations of modelling energy yield over the system's lifetime. If

the wind turbine cost, lifetime, and performance are assumed constant, minimizing the support structure cost will be equivalent to minimizing the COE, for a given siting scenario. Unfortunately, the dynamics of the system, if significant enough, can impact the turbine; larger motions potentially mean a more robust and expensive turbine is required, the turbine lifetime will be shorter, or additional control of the turbine to damp pitching motions in at-rated power operation will reduce the energy yield. Modelling these effects is outside the scope of the optimization framework. Instead, minimizing the platform motions that cause problematic turbine loadings is used as an optimization objective. The metric for the platform response/motions as affecting the wind turbine is the RMS fore-aft nacelle acceleration, equation (4.6). As mentioned in Chapter 4, this provides a good indication of the flapwise bending moments at the blade roots, which can be the most critical load in a wind turbine with a floating base.

The other factor affecting COE is the support structure cost. It can either be capped, or be capped and also minimized for in a multi-objective optimization[5].

For minimizing RMS nacelle acceleration, the objective function is of the form

$$\min J = \sum_{i=1}^{n} w_i (\sigma_{a\,nac.})_i \tag{5.12}$$

where w_i describes the weighting between the n metocean conditions being considered, with $\sum w = 1$, and $\sigma_{a\,nac.}$ is calculated according to (4.6). For a multi-objective optimization for RMS nacelle acceleration and cost, the objective function is of the form

$$\min J = W_1 J_1 + W_2 J_2 \tag{5.13}$$

where J_1 and J_2 are normalized RMS nacelle acceleration and cost, respectively, and W_1 and W_2 are weighting factors between the two objectives which sum to one. The normalization denominators used are 0.1 m/s^2 for RMS nacelle acceleration and \$10M for support structure cost.

Both approaches are employed in the results of Section 5.5. By considering turbine motion and support structure cost, the support structure design factors that are relevant for the COE of a floating offshore wind turbine are accounted for.

[5]Unlike cost, which is calculated before the hydrodynamic analysis, the platform motion is always treated as an objective rather than a constraint because there is no standard threshold for what amount of motion is acceptable and its calculation requires full evaluation of the design in the frequency-domain model.

5.4.3 Constraints

In addition to the basic geometric constraints included in the design parameterization discussed in Section 5.2, a number of performance constraints are applied to ensure candidate designs are feasible.

Costs

Total support structure cost, as calculated according to the functions of Section 5.2.6, was capped at $9M.

Static Pitch Angle

A limit of 10° is placed on the static pitch angle of the platform. This is a widely-used limit for floating wind turbines [58, 51, 38]. Larger static pitch angles imply significantly off-axis inflow of air to the turbine rotor, which can result in poor performance and larger structural loadings.

The frequency-domain model of the combined system cannot model the static offset of the system; therefore the static pitch angle, $\bar{\xi}_5$, needs to be evaluated separately. The static pitch angle is a function of the platform volume and center of buoyancy, the turbine and platform masses and centers of masses, the water plane moment of inertia, and the mooring system stiffness:

$$\bar{\xi}_5 = \frac{F_5}{C_{5,5}} = \frac{F_{thrust} z_{hub} + F_{l\,5} - \frac{1}{2} R_F N_F m_{ballast\,O} g}{\rho \forall g z_{CB} - m_t g z_{t\,CM} - m_p g z_{p\,CM} + \rho g I_{wp} - C_{l\,5,5} + C_{l\,5,1} z_{fair}} < 10^\circ \tag{5.14}$$

where F_5 is the total pitch moment on the system and $C_{5,5}$ is the total pitch stiffness. F_{thrust} is the thrust loading on the turbine, z_{hub} is the hub height (assumed to be on the line of the thrust force), $F_{l\,5}$ is the force in the pitch DOF exerted by the mooring system at the maximum-thrust equilibrium surge displacement, $\forall$ is platform displacement, z_{CB} is the center of buoyancy height, m_t is wind turbine mass, $z_{t\,CM}$ is turbine center of mass location, m_p is platform mass, $z_{p\,CM}$ is platform center of mass location, I_{wp} is platform water plane moment of inertia in the pitch direction, $C_{l\,5,5}$ is the stiffness in pitch from the mooring lines, and $C_{l\,5,1} z_{fair}$ is the product of pitch-surge mooring stiffness and fairlead depth. The last term accounts for the fact that the static pitch should occur about an axis located at the fairlead draft rather

than at the water plane; this avoids unrealistic overprediction of the pitch stiffness contribution of the mooring lines.

This constraint is evaluated at the maximum wind thrust condition. For the NREL 5 MW reference turbine, this is 800 kN, and the hub height is 90 m [24]. The mooring system properties – $F_{l\,5}$, $C_{l\,5,5}$, and $C_{l\,5,1} z_{fair}$ – are based on the corresponding surge displacement. Equation (5.14), like the rest of the model, assumes small angles.

Dynamic Pitch Angle

A dynamic pitch constraint is necessary to ensure operating angle limits for the turbine and floating platform are not exceeded. Following the approach of Tracy [51], a maximum steady plus RMS pitch angle of 10 degrees is used:

$$\bar{\xi}_5 + \sigma_{\xi_5} < 10^{\circ}. \tag{5.15}$$

The standard deviation in pitch, σ_{ξ_5}, is calculated based on the wave spectrum and the platform's pitch RAO as described in Section 2.3.4.

Slackness in Mooring Lines

Snap loads, where a taut mooring line goes slack and then abruptly regains tension, can cause large loads and structural failure for taut-moored support structures. Avoiding taut lines going slack is therefore a design constraint. Using the frequency-domain approach, the potential for the mooring lines going slack is calculated using RAOs for the mooring lines using the same RMS approach as for the pitching motions.

$$\bar{T}_{line} - 3\sigma_{T_{line}} > 0 \tag{5.16}$$

where $\bar{T}_{line}$ is the steady-state line tension and $\sigma_{T_{line}}$ is the RMS line tension variation about the mean calculated from the line tension RAO according to the approach as described in Section 2.3.4.

5.4.4 Inputs

The framework takes a number of inputs that characterize the operating environment of the floating wind turbine. These inputs are: water depth, a set of wave spectra, and a set of corresponding steady wind speeds. The site conditions would also have

implications for the costs associated with the structure, and anchor costs in particular. In addition to these site-specific inputs, a number of inputs relating to design assumptions and constraints are used. As has been discussed previously, these include the type of ballast to be used, the expected hull thicknesses, the thicknesses of mooring cables, upper limits on the structural mass and mooring line tensions to reflect cost constraints, and maximum acceptable pitch angles.

All of these inputs reflect the nature of the framework as a global optimizer; once an operating environment and common design constraints are provided, the framework will explore all the options within those inputs according to its abilities.

The site-specific input variables used for the results presented here are a water depth of 300 m, wind speeds of 8 m/s and 12 m/s, and corresponding sea states of 5 m and 8 m significant wave heights, and 6 s and 10 s peak periods. The two wind speeds correspond to when the turbine is in region II (constant pitch operation) and just after it has entered region III (power limiting operation) so that the thrust load is near maximum. The sea states are chosen to represent reasonable hypothetical wave conditions given these wind speeds. The two environmental conditions are given equal weighting; the objective function is the average of the objective function value calculated for each of the two environmental conditions.

The frequency range used in the analysis is from 0.25 rad/s to 2 rad/s, in 0.125 rad/s increments. The bottom of this range is below the wave excitation spectrum and the top of this range is above the active frequencies in typical RAOs.

5.4.5 Design Evaluation Implementation

When the genetic algorithm produces a point in the design space to be evaluated, a collection of subroutines in Matlab and a C DLL implement the processes described in Sections 5.2 and 5.3. The flow of these subroutines is shown in Figure 5.8.

5.5 Results

This section provides optimization results generated by the framework using the input conditions described previously, in order to demonstrate the framework's operation and present findings on the design space described by the framework's parameterization.

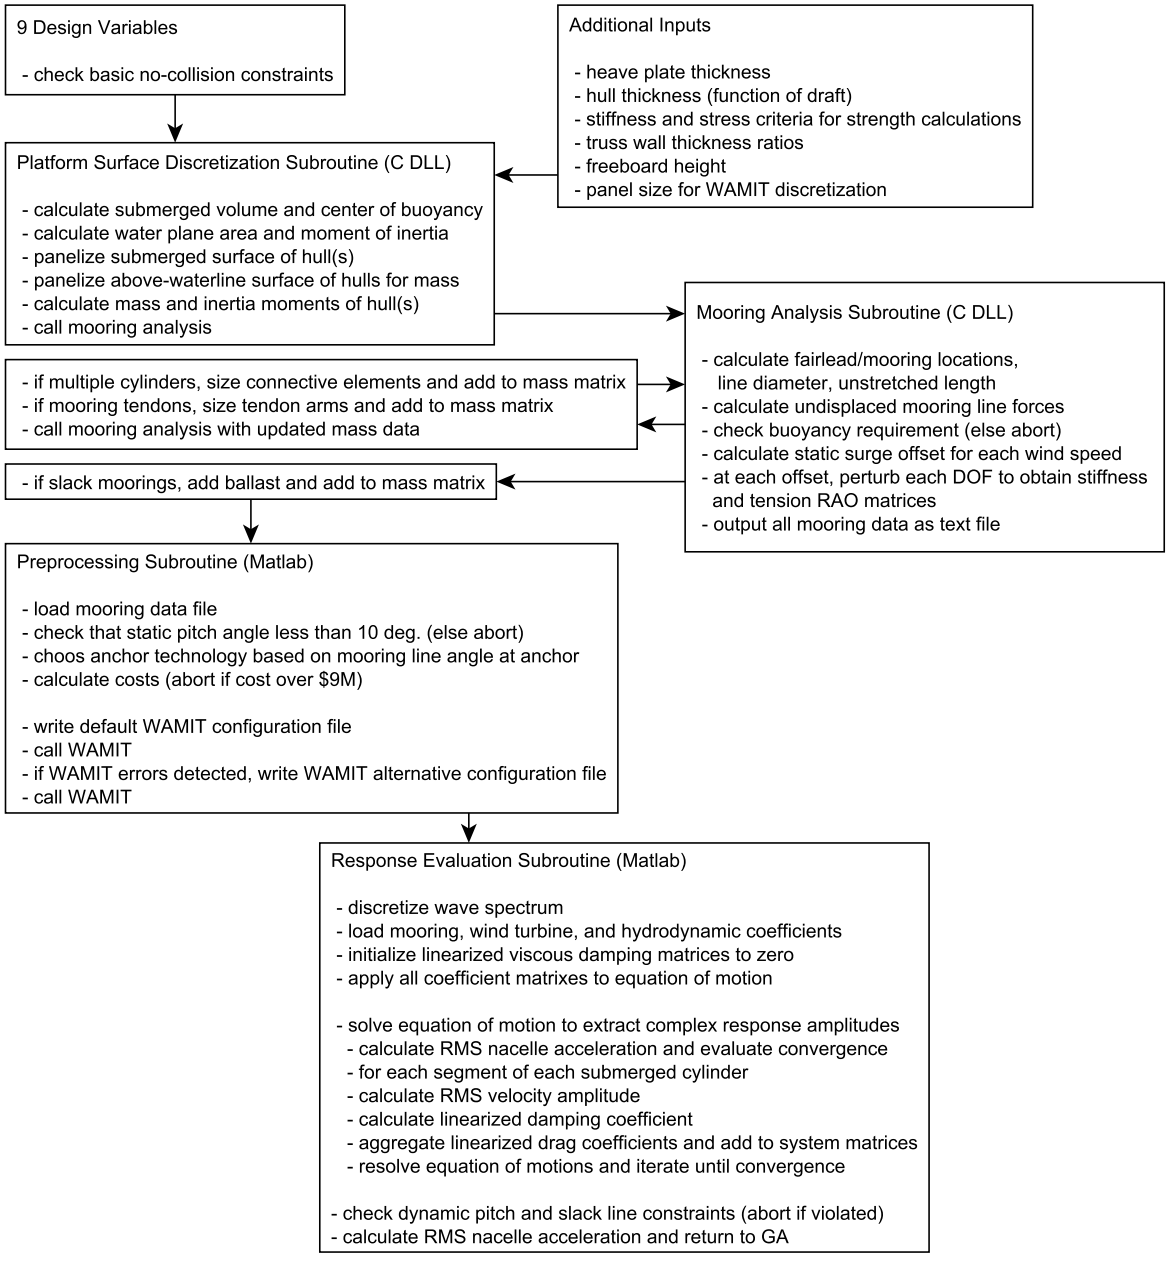

Figure 5.8: Flow diagram of design evaluation implementation

5.5.1 Single-Cylinder Single-Objective Optimization

One of the simplest demonstrations of the framework's operation can be made by considering a single-cylinder design space. The four variables describing this design space are draft, H_I, radius, R_I, taper ratio, T_I, and mooring configuration, x_M. The framework was run on this design space two times, with each run terminating at a population size of 1500. The computation time for such a run is in the order of one day on a four-core personal computer.

The three-dimensional scatter plot of Figure 5.9 shows the framework's exploration of this design space in terms of H_I, R_I, and x_M, for the first run. Taper ratio was found to add little complexity to the design space; whatever taper ratio gives the smallest feasible water plane area is generally the best. The colour of each data point represents that design's scaled fitness value as used by the GA. A value of zero indicates a least-fit individual and a value of one indicates a locally-optimal individual. The points evaluated by the GA can be seen to cluster around three configurations:

- spar-buoys on the right,
- wide ballasted cylinders on the left, and
- TLPs of various dimensions at the bottom.

This demonstrates the framework's ability to converge to multiple local optima, and shows that there are multiple local optima even in just the single-cylinder design space. The five locally-optimal design points identified in Figure 5.9 are illustrated in Figure 5.10. These points are only locally optimal in terms of the population of points evaluated by the framework.

The design variable values of the top five local optima from both the first and second runs are given in Table 5.3. Looking at the table, the optimal platform results of both cases are nearly identical for the first two local optima, but then different for the next three local optima, suggesting that the GA's results for the global optimum are reliable.

The spar buoy design has a draft of 132 m, which compares closely to the 120 m draft of the Hywind design. The diameter is greater than the Hywind design (11.8 m compared to 9.4 m) – not surprising since the GA will tend to seek the largest structure that meets the \$9M cost constraint, whereas the Hywind would have been designed with cost minimization in mind.

Table 5.3: Single cylinder results comparison

	Run 1 Optima no.					Run 2 Optimum no.				
	1	2	3	4	5	1	2	3	4	5
H_I (m)	132.2	26.2	53.9	106.1	17.7	123.9	26.6	133.9	9.3	19.7
R_I (m)	5.92	23.09	3.64	3.30	4.59	6.45	22.30	3.42	6.68	4.33
T_I	0.51	0.49	0.85	0.92	1.53	0.47	0.42	0.89	1.29	1.70
x_M	1.87	1.89	-0.98	-0.49	-0.36	1.78	1.85	-0.36	-0.47	-0.85
$\sigma_{a\,nac.}$ (m/s^2)	0.139	0.172	0.270	0.285	0.367	0.141	0.177	0.261	0.348	0.348
cost (k$)	8888	8669	8210	7306	5040	8953	8355	8835	5732	8774
$m_{plat.}$ (tonnes)	11879	37290	2548	2141	1137	13196	34662	2573	1546	2805
$\forall$ (m^3)	12451	37375	4474	4363	2336	13739	34796	5335	2851	4512
$T_{fair.max}$ (kN)	1177	1551	3663	4233	4128	1404	1763	5474	3474	3332

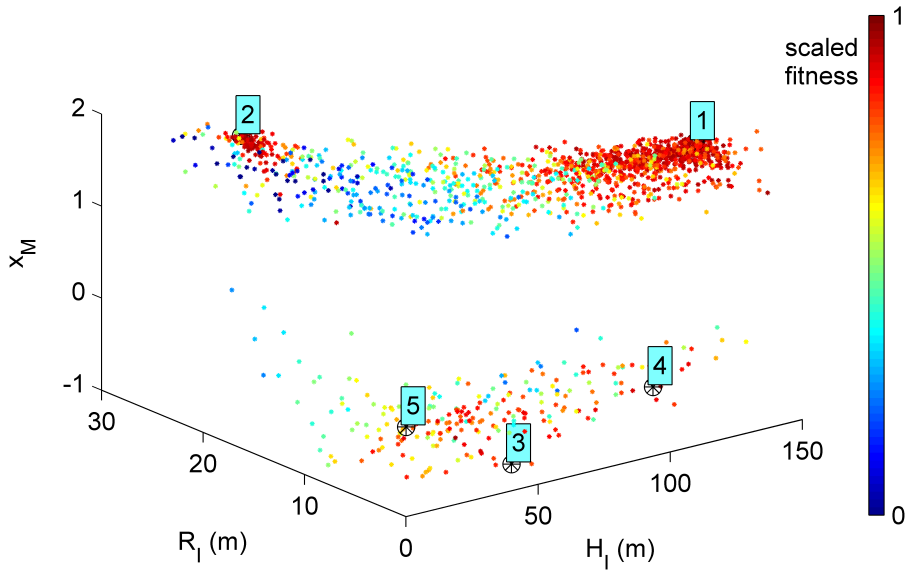

Figure 5.9: Single-cylinder single-objective design space exploration

The second locally-optimal design is a large cylinder with 31 700 tonnes of ballast (compared to 9 900 tonnes for the spar-buoy), resulting in triple the displaced volume of the spar-buoy. This design relies on a large amount of ballast and sheer size to stabilize against wind and wave forces. While it represents an interesting local optimum in the design space, the fact that it weighs over three times as much as the next-heaviest design makes it an unlikely candidate for real-world application. A cost for ballast is excluded from the framework, but adding a reasonable ballast cost does not eliminate this design. The design may face feasibility issues due to manufacturing, transportation, and structural issues associated with its size. With its significant amount of volume submerged less than 6 m below the waterline, this design likely also exploits the amplitude-independence assumptions of linear hydrodynamics. Realistically, in large-amplitude waves, portions of the submerged volume could become exposed, resulting in dangerous sudden load changes on the structure.

The third locally-optimal design is a TLP with a 54 m draft and 3.6 m radius, with tendon arms of similar dimension. This design seeks its stability from a balance between slender hull and widely-spaced mooring lines.

The fourth locally-optimal design is another TLP, this time with a 106 m draft and shorter tendon arms.

The fifth locally-optimal design is a much shorter TLP, with a draft of just 18 m

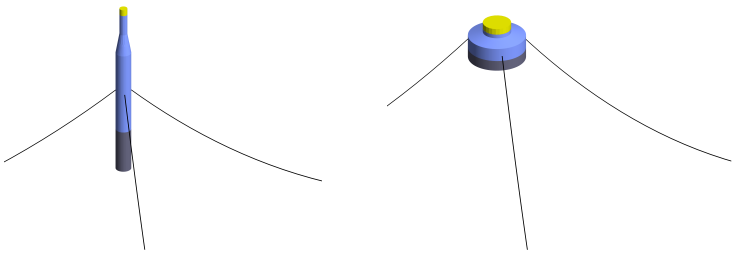

(a) Local Optimum 1 (b) Optimum 2

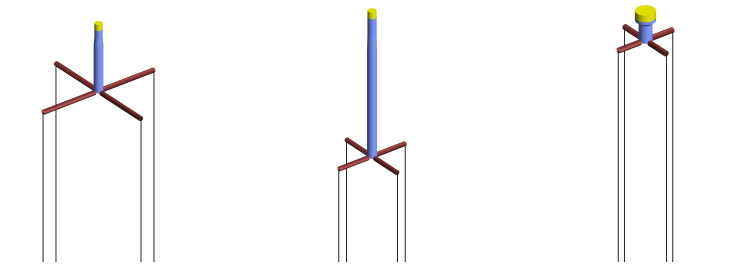

(c) Optimum 3 (d) Optimum 4 (e) Optimum 5

Figure 5.10: Single cylinder local optima

and a wide hull radius of 4.6 m at draft and 7 m at the water plane. Its greater water plane area exposes this design to more wave forces, causing its large nacelle acceleration value, but it's low surface area to volume ratio makes this design significantly less expensive than the others.

Between the two optimization runs, the third, fourth, and fifth locally-optimal designs are different, though they are all TLPs in both cases. There is evidence that both runs, if given more time, would converge toward the same local optima. For instance, from looking at Table 5.3 it can be seen that Optimum 4 in run 1 is similar to Optimum 3 in run 2.

Figure 5.11 shows a scatter plot of the first single-objective optimization population, with axes of support structure cost and RMS nacelle acceleration. The designs in the figure are classified as either:

- "TLP" for taut mooring systems,
- "Semisub" for multi-cylinder designs with slack mooring systems, or
- "Spar-buoy" for single-cylinder designs with slack mooring systems.

Plotting in terms of the two optimization objectives allows for visualization of

the performance space, the range of performance possibilities for feasible designs. From Figure 5.11, there is evidence of a boundary along the lower-left extent of the population. This boundary, called a Pareto front, represents the best performance possibilities given an unknown weighting between the two design objectives. Points on the Pareto front represent nondominated designs – designs that cannot be changed to achieve an improvement in one objective without worsening the other objective. Apparently, the nature of the GA is such that the rough shape of a Pareto front can be found in the single-cylinder performance space without needing a multi-objective optimization formulation.

It should be noted that with the stochastic nature of a GA and the practical limitations imposed on the population size, the extent of the population can only approximate the location of the Pareto front; designs produced by the GA can get close to the Pareto front limit, but will never fall exactly on it. Precautions have been taken to ensure that the population has come reasonably close to the Pareto front. This was done for the results in this section and the following sections by running additional optimizations – some with larger population limits and some with higher mutation rates – and confirming that none of the designs in these optimizations' populations out-performed the nondominated designs in the original optimizations by a significant amount.

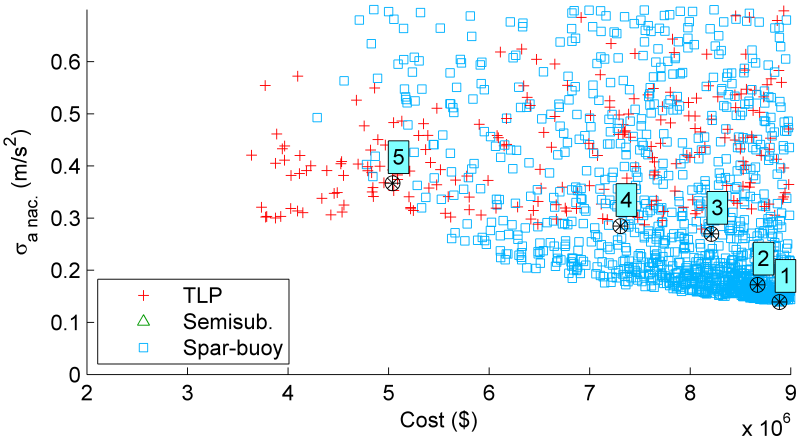

Figure 5.11: Single-cylinder single-objective performance space

5.5.2 Single-Cylinder Multi-Objective Optimization

Changing from a single-objective nacelle-acceleration optimization to a weighted sum of both nacelle acceleration and cost and selecting appropriate weightings causes the GA to converge toward lower-cost designs. Superimposing the populations from three simulations with different weightings (detailed in Table 5.4) results in a more expansive exploration of the design space (Figure 5.12) and better resolution of a Pareto front (Figure 5.13). Each multi-objective optimization run was terminated once the design population reached 1000.

Table 5.4: Weightings for singly-cylinder multi-objective optimization runs

	Run 1	Run 2	Run 3
W_1	1	0.5	0.09
W_2	0	0.5	0.91

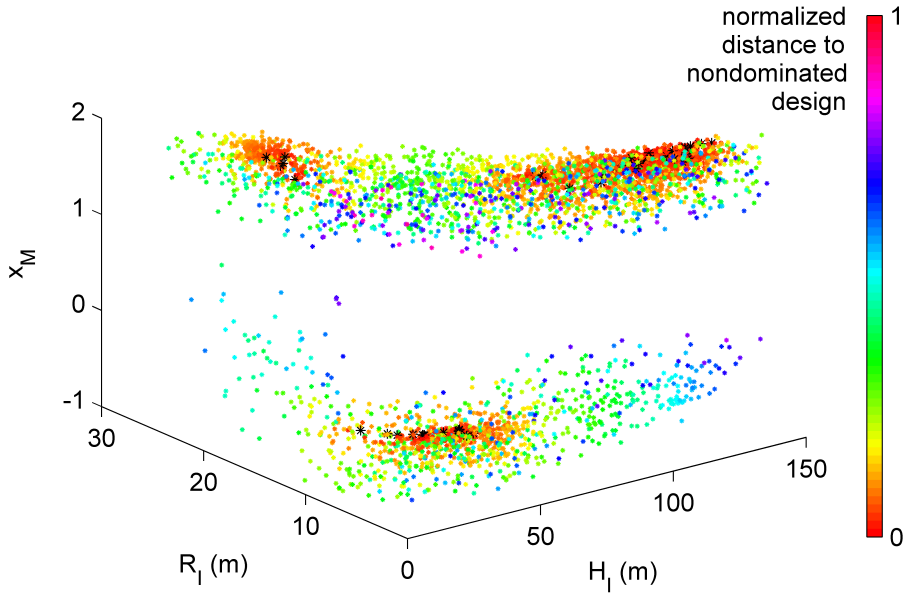

Figure 5.12: Single-cylinder multi-objective design space explorations

As can be seen, the spar buoy configuration is the most stable above a cost of about \$5M. Within that range, the large ballasted cylinder designs are competitive with more conventional spar-buoys for costs around \$5M to \$7M. Below \$5M, a TLP configuration can achieve greater stability for a given cost but cannot achieve as low accelerations as the higher-cost spar-buoys. No buoyancy-stabilized (barge-type) platform designs are to be seen on the Pareto front.

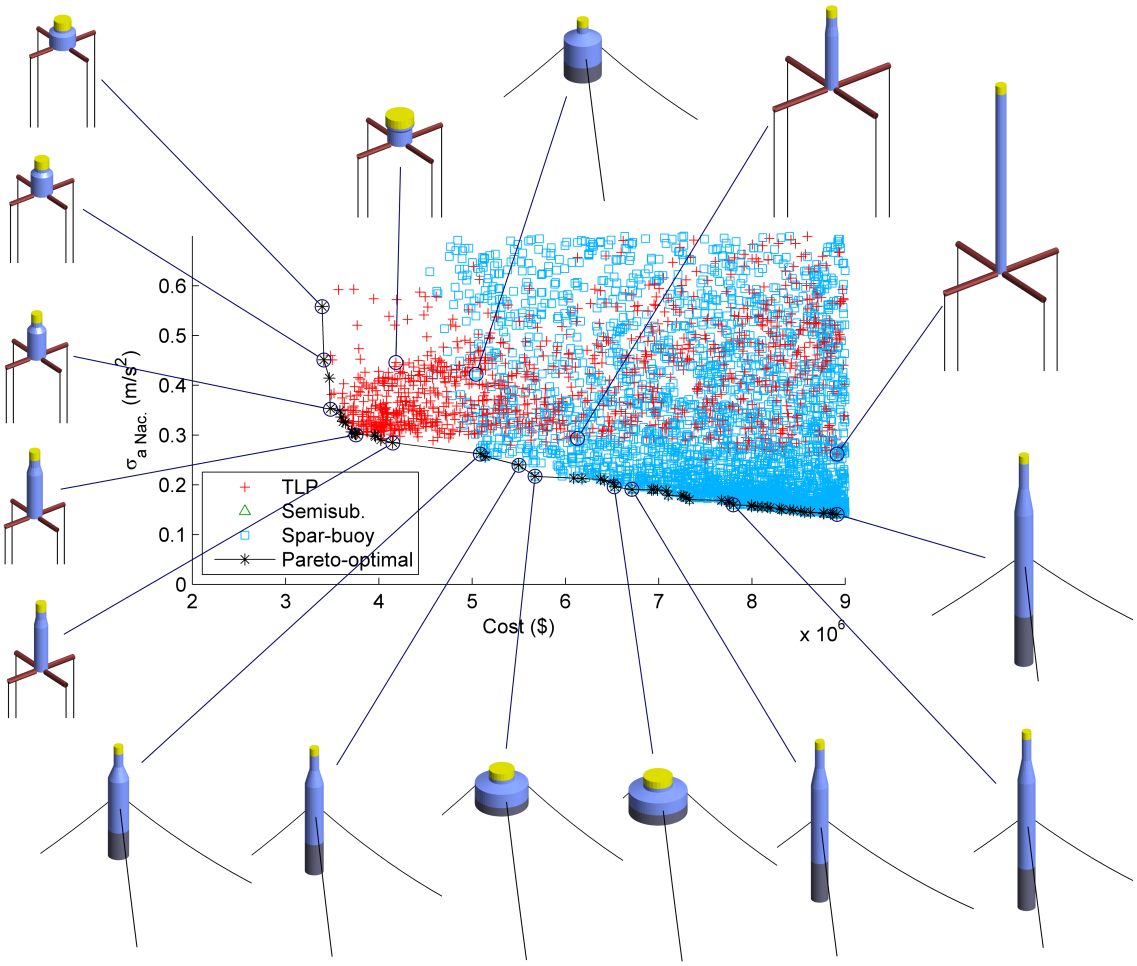

Figure 5.13: Single-cylinder multi-objective performance space

The platform illustrations in Figure 5.13 show how the configuration of the Pareto-optimal design changes with platform cost. There is a consistent trend along the entire Pareto front of deeper and narrower platforms as the cost increases. There is little variation in the moooring system, with $x_M \in [1.8, 1.95]$ above \$5M and $x_M \in [-0.4, -0.3]$ below \$5M. Based on the costs, constraints, and environmental conditions specified for these results, it seems that there is little chance of feasible support structures costing less than \$3.4M.

5.5.3 Full Design Space Single-Objective Optimization

When applied over the full design space described by the nine parameters discussed in Section 5.2, the framework identifies a number of locally-optimal designs. A visualization of the framework's exploration, in a three-dimensional projection of the design space with axes for R_I, R_O, and x_M, is shown in Figure 5.14. This result shows an optimization for nacelle acceleration only, with cost capped at \$9M, and a population size of 1500. The computation time for these runs is about twice that of the single-cylinder runs because of the larger number of panels for the WAMIT analysis of multi-cylinder platforms.

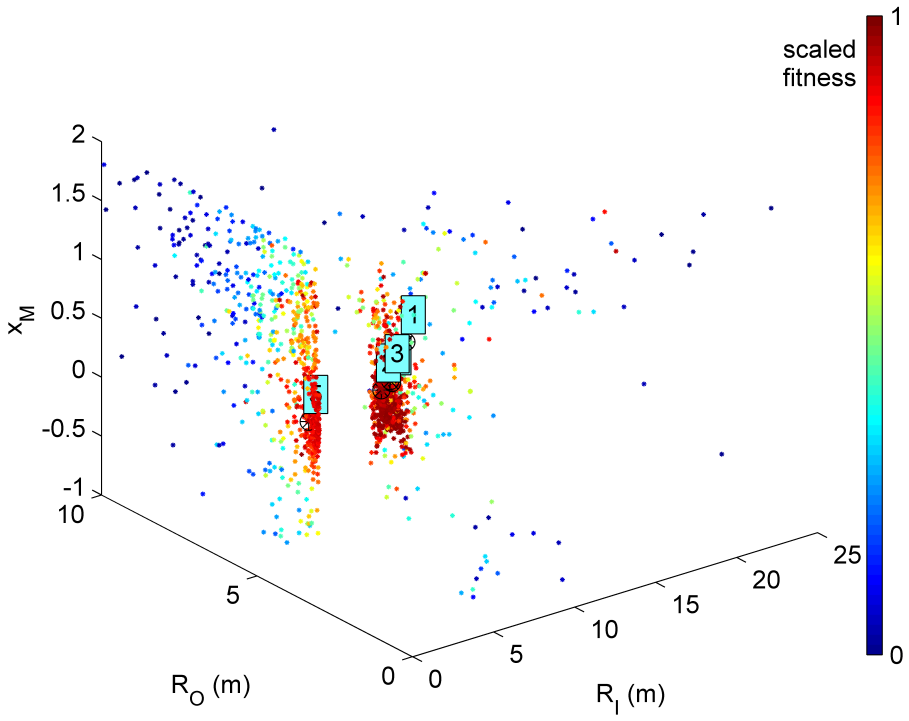

Figure 5.14: Full single-objective design space exploration

In the projection of Figure 5.14, two clusters of designs are visible. One cluster is on the $R_I = 0$ plane, indicating designs that do not have a central cylinder. It should be remembered that the evaluated designs are not as similar as they look from the figure, because the other six design space dimensions are collapsed in this plot.

The top five locally-optimal designs are illustrated in Figure 5.15. Details of these designs are provided in Table 5.5.

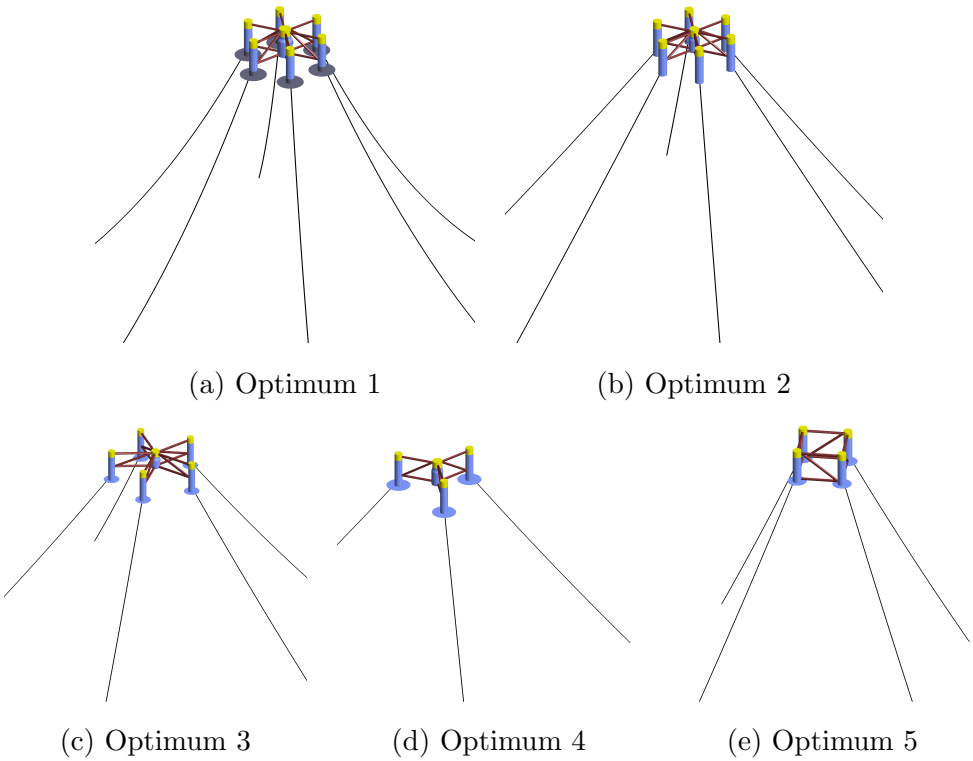

(a) Optimum 1 (b) Optimum 2

(c) Optimum 3 (d) Optimum 4 (e) Optimum 5

Figure 5.15: Full design space local optima

Local optimum 1 is slack-moored semisubmersible design with a small central cylinder and six cylinders arrayed around it. These outer cylinders have sizeable heave plates. Local optimum 2 features a similar hull arrangement but has no ballast or heave plates and uses a taut mooring system with a 45° line angle. Local optimum 3 is again similar but now with five outer cylinders, and heave plates in addition to the taut mooring system. Local optimum 4 reduces to three outer cylinders and larger heave plates. Local optimum 5 continues the trend of a taut mooring system with non-vertical lines. It features four cylinders arranged in a square, equipped with heave plates.

Taking the population of designs shown in Figure 5.14 and plotting them in Figure 5.16 in terms of cost and nacelle acceleration shows only a hint of a Pareto front. A dedicated approach for mapping the Pareto front is therefore required.

Table 5.5: Full design space local optima

Optima no.	1	2	3	4	5
N_F	6	6	5	3	4
R_F (m)	26.8	26.7	37.0	31.0	24.1
H_I (m)	15.3	12.2	10.9	14.9	0.0
R_I (m)	5.18	3.81	4.12	4.40	0.00
T_I	0.783	0.852	0.862	0.857	0.200
H_O (m)	21.5	21.3	22.4	22.3	21.5
R_O (m)	2.90	2.99	2.87	2.95	3.34
R_{HP} (m)	9.1	0.0	6.6	9.0	7.3
x_M	1.06	0.69	0.78	0.73	0.54
$\sigma_{a\,nac.}$ (m/s^2)	0.100	0.153	0.160	0.164	0.172
cost (k$)	7374	8532	8129	5282	6610
$C_{platform}$ (k$)	6538	5625	6786	4024	5243
C_{lines} (k$)	524	1436	684	630	653
$C_{anchors}$ (k$)	312	1471	658	628	713
$m_{platform}$ (tonnes)	4004	2250	2714	1610	2097
$\forall_{platform}$ (m^3)	4805	4342	3927	2857	3354
$T_{line\,max}$ (kN)	935	3323	1808	2836	2428

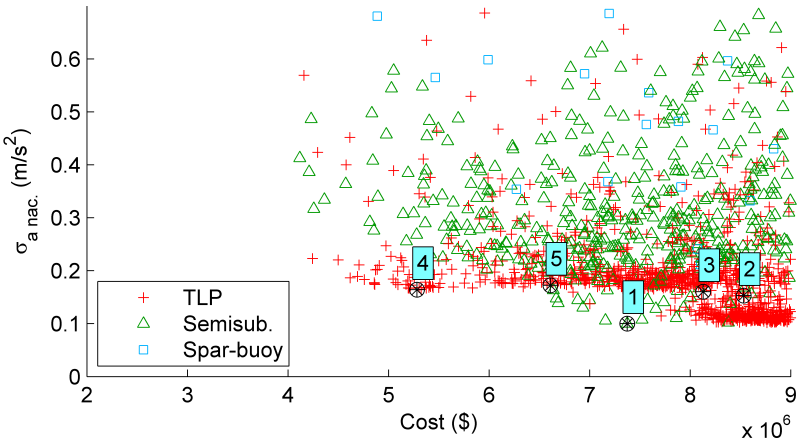

Figure 5.16: Full single-objective performance space

5.5.4 Full Design Space Multi-Objective Optimization

Resolving a Pareto front for the full design space is considerably more involved than for the single-cylinder design space, necessitating more weightings of the objective functions. This is because the possibilities for different numbers of cylinders and different mooring systems create many different niches for the GA to explore. These niches inhibit the GA from exploring a wide span of the Pareto front in a single optimization run of a given weighting.

Seven weightings (detailed in Table 5.6) are used to generate the population shown in the design space by Figure 5.17 and in the performance space by Figure 5.18. Black points in these figures indicate nondominated individuals. A population size of 1500 was used in the optimization for each weighting.

Table 5.6: Weightings for full design space multi-objective optimization runs

	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7
W_1	1	0.833	0.667	0.286	0.167	0.0625	0.032
W_2	0	0.167	0.333	0.714	0.833	0.9375	0.968

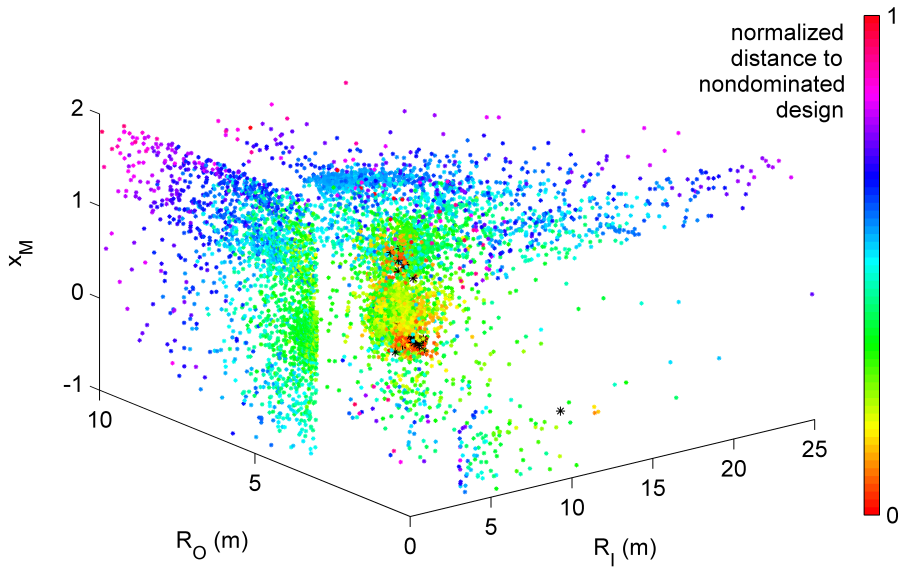

Figure 5.17: Full multi-objective design space explorations

Though the single-objective optimization showed some variety in the local optima discovered by the framework, the multi-objective optimization shows only two general platform configuration types along the Pareto front.

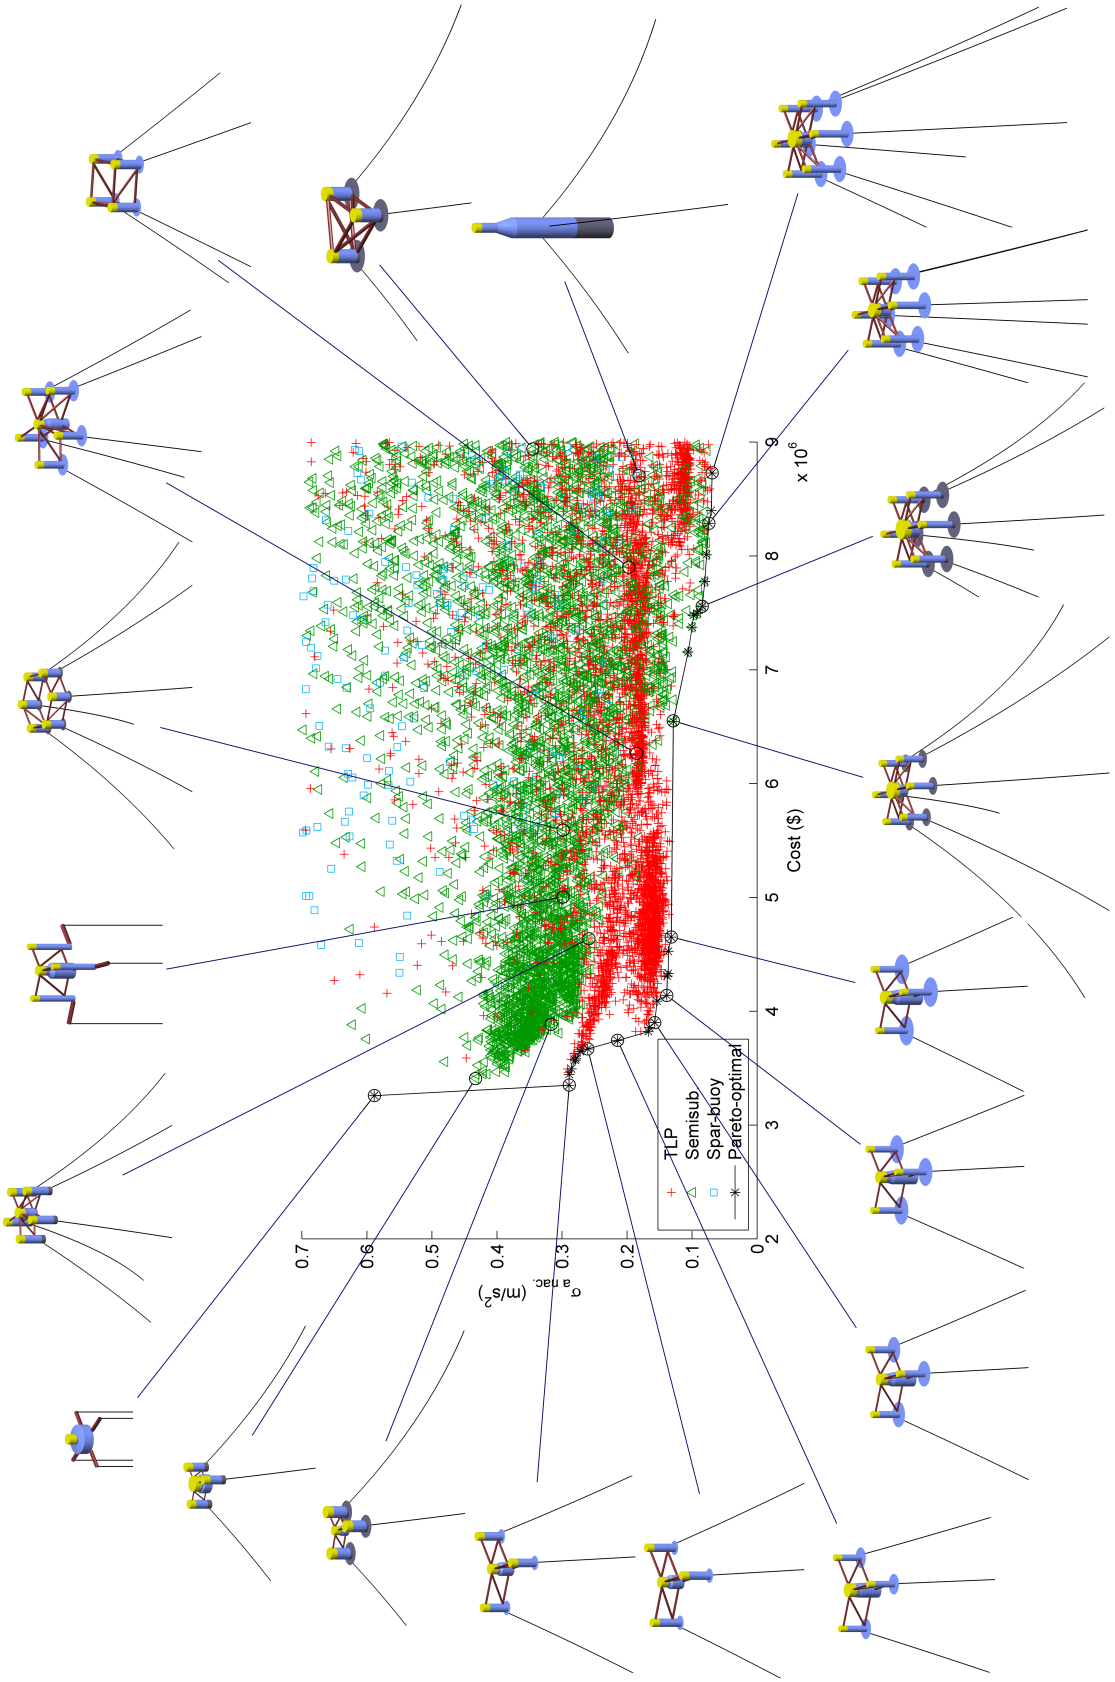

Figure 5.18: Full multi-objective performance space

Above a cost of $6M, the Pareto-optimal platform configurations feature six slender outer cylinders arrayed around a shorter central cylinder. Heave plates are used in all cases. The mooring system transitions from a taut system with non-vertical lines to a slack system as cost decreases, with ballast added to compensate. Heave plate size reduces as cost decreases.

Below a cost of $6M, the Pareto-optimal platform configurations feature three slender cylinders arrayed around a central cylinder of similar draft but larger radius. The mooring system is taut but non-vertical, and the cylinder spacing can be seen to increase slightly, while the heave plate size decreases, as cost decreases. It seems that cylinder spacing is more economical but less effective than heave plate area at reducing nacelle acceleration.

The lowest-cost non-dominated design departs from the others in being a single-cylinder shallow-draft TLP design. Its wide platform shape provides a very low surface area to volume ratio. It would seem that mooring line tension is the most economical way of stabilizing such a platform. Based on the costs, constraints, and environmental conditions specified for these results, it seems that there is little chance of feasible support structures costing less than $3.4M.

5.5.5 Time-Domain Verification of Global Optimum

Recreating some of the designs produced in the framework and evaluating them in the higher-fidelity time-domain FAST simulation is one way to check the validity of the framework's results. This was done with the spar-buoy configuration that was found to be the global optimum in the single-cylinder optimization, pictured in Figure 5.10(a). This design is an advantageous choice because it avoids the truss elements and heave plates of other designs, whose viscous hydrodynamic properties FAST is unable to account for. A comparison to FAST time-domain results with a more complex platform configuration such as that of Figure 5.15(a) would be less informative, because FAST would neglect the significant viscous drag forces on the connective members and heave plates.

The mooring system properties of the design, before linearization, can be used directly in the FAST platform input file, as can the platform mass properties. The hydrodynamic coefficient data files generated by WAMIT in the framework can also be used directly with FAST. To test the adequacy of the meshing, a finer discretization of the platform geometry was also created (with 8204 panels compared to the original 2076). The original and finer-discretization platform geometry files used to generate the WAMIT data for the time-domain simulation are shown in Figure 5.19.

(a) Default discretization (b) Fine discretization

Figure 5.19: Spar-buoy surface meshes input to WAMIT

The steady winds used in the framework are easily implemented with the full time-domain aero-elastic model in FAST. The irregular sea states used in the framework are recreated in the time domain in FAST. Additional stiffness in the yaw direction

was necessary to maintain a stable turbine heading[6]. A separate simulation was done for each of the two metocean conditions used in the framework. Each simulation was run for 20 minutes, with the first 5 minutes of data excluded from the analysis to avoid start-up transients.

The calculated RMS fore-aft nacelle acceleration and RMS mooring line 3 tension are compared for the two FAST simulations and the framework's frequency-domain analysis across the two metocean conditions in Table 5.7. Mooring line 3 is one of the windward mooring lines, making it a good subject for mooring tension study. Agreement is very close between the FAST simulations with the original and the finer platform surface discretizations, indicating that the discretization used in the framework is adequate.

Comparing the results of the framework's frequency-domain model and the FAST time-domain simulations, there is significant disagreement. The RMS nacelle acceleration is less in the framework's model than in FAST by about 25% for the milder metocean conditions and 15% for the stronger metocean conditions. This is not surprising, considering that the framework's model does not account for the flexibility in the tower, which would increase the nacelle accelerations. The RMS tension fluctuations in the third mooring line disagree by more than an order of magnitude. This is likely an issue with the frequency-domain approximation of superimposing the fluctuations at each frequency, though the exact cause has not been determined. The steady or average mooring line tensions agree reasonably well between the two approaches, as is to be expected considering they both rely on the same quasi-static mooring model for determining the static surge offset of the platform. Overprediction by the framework's model relative to the FAST simulation of about 5% in both metocean conditions can be explained by the framework model's assumption of zero static pitch angle during the mooring line linearization. In the FAST simulation, a steady pitch angle of between 2° and 4° causes a slight reduction in rotor thrust which in turn reduces the tension on the windward lines.

Providing a more detailed check on the validity of the framework's frequency-

[6]The small radius of a spar-buoy floating wind turbine necessitates a "crow foot" mooring line attachment in order to provide sufficient stiffness in the yaw DOF. The delta connection in the mooring lines is not supported by conventional quasi-static models; it is therefore common practice to model the mooring system without the crow foot and add artificial yaw stiffness to the simulation to compensate [23]. For the present simulation, an amount equal to the amount used in the OC3 Hywind design model was used. This additional stiffness is not needed in a frequency-domain model because periodic yaw motions are very small and gradual yaw drift is precluded by the periodic assumptions.

Table 5.7: Comparison of frequency- and time-domain results

	8 m/s wind speed			12 m/s wind speed		
	F'work	FAST	FAST hi-res	F'work	FAST	FAST hi-res
$\sigma_{a\,nac.}$ (m/s^2)	0.089	0.121	0.123	0.189	0.223	0.225
$\sigma_{T\,line3}$ (kN)	0.64	44.5	43.4	8.52	89.0	87.0
$\mu_{T\,line3}$ (kN)	779.6	779.5	778.8	946.8	932.8	931.0

domain model is difficult because there is a lack of adequate sources of comparison. Scale physical model tests are very few and the available data is very limited. As well, scaling issues in these experiments limit their validity. On the computer modelling side, current time-domain models lack adequate treatment of the viscous drag effects which the present framework's model seeks to include, and there are large discrepancies between the best experimental results and numerical modelling (such as in [25]). The well-established nature of the individual models that have been combined to make the overall frequency-domain model is therefore relied on to support the validity of the calculated results involving viscous drag forces. To provide a cursory verification of the overall model,the RAOs of two designs in the framework are compared to the published figures for the OC3 Hywind and OC4 WindFloat designs in Appendix C.

5.6 Conclusions

A global optimization framework has been developed for the floating wind turbine support structure design problem. A platform geometry scheme based on arrays of vertical cylinders and a mooring configuration scheme with one dedicated decision variable provide a flexible and efficient means of describing a wide range of support structure configurations. A frequency-domain model evaluates the support structure performance in terms of platform motions in six degrees of freedom. The panel method code WAMIT, supplemented by a custom viscous drag model, provides platform hydrodynamic characteristics, while a quasi-static mooring model provides linearized mooring system stiffnesses and the floating wind turbine simulator FAST provides linearized wind turbine effects. These inputs are coupled together and solved as a function of frequency to evaluate the platform motions. A genetic algorithm controls the exploration of the design space, seeking local optima. The primary objective used in the optimization is the minimization of RMS nacelle acceleration. Constraints are implemented to represent a variety of performance requirements. Cost can be

treated as a constraint, or as a second objective function in a weighted multi-objective optimization. Nacelle acceleration and support structure cost together constitute the most relevant support structure design factors affecting the cost of energy from a floating wind turbine.

Results produced from the framework using hypothetical input data demonstrate the capabilities of the framework and reveal various characteristics of the design space. The framework converges reliably to locally-optimal designs. Some of these designs are very similar to existing designs; the best example is the single-cylinder global optimum pictured in Figure 5.15(a), which is similar to the Hywind design. This is valuable confirmation of consistency between the framework's models and real-world design factors. In the full design space, however, the framework finds that less common configurations featuring multiple cylinders perform better. These designs feature a central cylinder surrounded by three or six outer cylinders equipped with heave plates, and often have taut non-vertical mooring lines.

Surrounding these locally-optimal points in the design space are large swaths of feasible space. The GA approach of the framework is able to map these regions, allowing visualization of the nature of the design space – including the bounds imposed by expense, buoyancy, or stability constraints; the general effects of different parameters on the support structure's performance; and the presence of multiple local optima.

By viewing the results in terms of both cost and stability objectives, a Pareto front can be observed. Clear trends are visible in the designs as one moves along the front: three outer cylinders are best below a cost of \$6M, six outer cylinders are best above a cost of \$6M, and heave plate size increases with support structure cost. With the current settings, there appears to be a floor on support structure cost at around \$3.4M. At no point does a spar-buoy design sit on the Pareto front.

The presence of these relatively complex four- and seven-cylinder platforms on the Pareto front, and the absence of simpler buoyancy- or ballast-stabilized designs, would appear to challenge the conventional wisdom in support structure design for floating wind turbines. However, the complexity of these designs may carry additional costs and risks not accounted for in the framework. If accounted for, these factors could tip the balance back toward more conventional single- or three-cylinder platform designs.

In addition to identifying non-dominated designs, the framework provides the ability to map the performance space, clarifying where the tradeoffs between different support structure classes occur. With more detailed input data, this framework could be used to help navigate design tradeoffs for specific site scenarios, determining in an

efficient, fair, methodical, and reliable fashion which configuration is best.

5.7 Future Work

Further development of the optimization framework could be carried out in several areas. One area for improvement is in expanding the cost model to better account for the additional costs associated with more complex platform designs.

Another area for improvement is in the framework's design evaluation process: improving the coupling between static pitch angle, wind turbine linearization point, and mooring system linearization point; and improving the sizing algorithms of structural elements and mooring lines to better reflect the dynamic loads they will face. Most of these changes necessitate a more iterative design evaluation approach, in which the dynamics results feed back into the sizing of the structure components until a well-sized design is converged upon. Good data on structural sizing, costs, site-specific conditions, etc. would be necessary to make this worthwhile. Until these input data are more trustworthy, the proposed iterative approach would be a significant increase in computational expense with still-inaccurate results.

The limitations of frequency-domain modelling for representing certain loads is also an important issue. Particularly in light of the disagreements demonstrated in Table 5.7, there would be value in developing improved methods of representing realistic loads in a frequency-domain model. Empirical approaches for approximating the damping contributions of different mooring configurations could also be considered. Alternatively, if the computational resources were available, higher-fidelity time-domain modelling could be used in the framework, to provide a marked increase in the quantity and quality of evaluation data available.

The support structure parameterization could be significantly expanded, to include a greater variety of design possibilities. One specific example of this would be to add support for designs with high tilt angles, such as the Sway design, by performing the hydrodynamic analysis on the tilted structure after the static pitch angle has been calculated, and engaging the nonlinear mooring model in the static pitch calculation. Increased intelligence in the use of advanced features such as active ballast and different mooring line options could further expand the design space. Accounting for different wind turbine positions, as is realistic for platforms without a central cylinder, could also be done. Additional constraints could be added as well, for example: accounting for the effect of tidal variations when evaluating whether a

taut-moored design will experience snap loadings or restricting the natural periods of the floating platform from a certain range (as is done in [56]). It is challenging to develop flexible parameterizations that maintain order in the design space, but there may be potential for alternative approaches that do not rely on cylindrical geometries.

Related to the parameterization is the question of how similarity or distance in the design space is calculated. The framework currently uses a simple Euclidean distance calculation with each of the decision variables, but this can in some cases cause problems. For example, difficulties were encountered wherein the GA produced an overpopulation of spar buoys with minuscule cylinders arrayed around them in varying numbers and at varying radii. Because the tiny size of the cylinders was not accounted for in the distance calculation, the algorithm produced many of these designs thinking they were dissimilar when in fact they were almost identical in performance. More intelligent ways of measuring similarity and distance in the design space would improve the optimization operation. Another example of this is in Figure 5.11 – good TLP designs on the left of the plot are not identified as local optima because of supposed similarity with local optima 3, 4, and 5. Many of these designs are in fact dissimilar. To reflect this, the similarity comparison could be based on the ratios between design variables rather than simply the Euclidean distance between the design variables.

Lastly, the GA itself could be improved to provide a more thorough mapping of the performance space. One possibility would be to adjust the structure of the algorithm so that it seeks Pareto-optimal rather than locally-optimal designs, possibly by making the distance threshold function be a function of an offspring's distance from the nearest non-dominated individual. If this is not possible, an alternative would be to leave the GA as-is but to add a technique to help ensure that, with enough optimization runs, even concave regions of the Pareto front are well mapped – the adaptive weighted-sum method could be used for this [69].

Chapter 6

Conclusions

This thesis presents three different explorations into some of the technical challenges associated with floating offshore wind turbines. The value of these explorations is not only in the conclusions they produced, but also in the directions for future work that they point toward.

6.1 Adequacy of Quasi-Static Mooring Models

To look into one of the lesser-explored modelling accuracy questions in floating wind turbines, mooring line modelling, a dynamic FEM-based mooring line model was coupled to the floating wind turbine simulator FAST and a broad comparison of the results produced by quasi-static and dynamic mooring models was done. This comparison yielded a number of conclusions:

- Damping from the mooring lines can make up a significant portion of the overall platform damping for some designs – particularly when platform motions are large and the fairlead locations are far from the center of mass.
- Quasi-static mooring models are well-suited for slack-moored designs that have natural frequencies well below the peak wave periods, small motions, and fairleads located close to the platform's center of mass (eg. spar-buoys).
- Quasi-static mooring models are not suited to slack-moored designs that have natural frequencies near or above the peak wave periods, large motions, and fairleads located far from the platform's center of mass (eg. barges).

- The suitability of quasi-static mooring models for slack-moored designs that fall in between the two above extremes in terms of natural frequencies, motion amplitudes, and fairlead locations should be evaluated on a case-by-case basis using a comparison study similar to this one.
- For taught-moored designs, quasi-static mooring models can provide an approximation of the system dynamics but cannot provide high accuracy in turbine load prediction.
- Small inaccuracies in the platform motion time series introduced by a quasi-static mooring model can cause much larger inaccuracies in the time series of the higher-frequency rotor blade dynamics.

6.2 Basis Function Platform Optimization

To explore a "hydrodynamics-based" approach to floating platform optimization, a basis function technique was developed in which the design space was represented by the linear combination of performance characteristics from a set of geometrically-defined "basis" platform designs. This exploration yielded conclusions about both the design space and the challenges of a hydrodynamics-based optimization paradigm:

- A basis-function optimization approach yields solutions favouring combinations of large submersed volume well below the waterline and widely-distributed water plane area, regardless of the mooring system being used. The preferred design is a combination of a submersed cylinder and a three-cylinder semisubmersible.
- The physical consistency of the basis function approach of linearly combining platform characteristics only holds for two extreme cases:
  - combinations of similarly-shaped geometries, which can represent an intermediate geometry, and
  - combinations of well-separated geometries, which can represent a geometry featuring scaled down versions of each basis geometry.
- Combinations of geometries that fall in between these interpretations, such as intersecting but differently-shaped platforms, are not well represented by linear combination because of hydrodynamic interactions between the platforms.

6.3 GA-Based Support Structure Optimization

To improve support structure design space exploration capabilities, a genetic algorithm-based support structure optimization framework was developed. This framework embodies contributions in design-space parameterization, automated frequency-domain modelling, and metaheuristic-based optimization techniques. Application of this framework to a hypothetical scenario confirmed the framework's effectiveness and yielded a number of conclusions about the design space:

- The framework converges reliably to locally-optimal designs, and identifies both conventional and unconventional support structure configurations.
- The framework is well-suited to exploring the design space and providing visualization of bounds and trends, including mapping Pareto fronts.
- The floating wind turbine support structure design space has multiple local optima.
- The floating wind turbine performance space as defined by support structure cost and RMS nacelle acceleration has a clear Pareto front. According to the models used in the framework:
  - a seven-cylinder platform with heave plates gives the least nacelle acceleration for support structure costs above $6M,
  - a four-cylinder platform with heave plates and taut non-vertical mooring lines gives the least nacelle acceleration for support structures costs below $6M,
  - a floor on support structure cost exists near $3.4M,
  - conventional spar-buoy and TLP configurations are dominated by multi-cylinder configurations; however, their relative simplicity over the non-dominated designs found by the framework means that single-cylinder platforms should not be ruled out.

6.4 Future Work

6.4.1 Adequacy of Quasi-Static Mooring Models

The work of Chapter 3 was fairly self-contained in presenting a methodology for evaluating the adequacy of quasi-static mooring line models, demonstrating that the results are highly dependent on the support structure configuration, and providing some general guidelines for the situations in which quasi-static mooring models are adequate. There are, however, clear directions for further research into mooring line modelling for floating wind turbines:

- Finite-element mooring line models are very computationally demanding, and bending and torsional stiffnesses were found to have a negligible effect on system dynamics results. There is therefore good reason to explore the suitability of simpler lumped-mass dynamic mooring line models that neglect bending stiffness for floating wind turbine simulation in order to retain the dynamic effects of the mooring lines without the computational cost of a finite-element model.

- Only two fully-stochastic operating conditions were used in the present study, and the disagreement between mooring models was found to be sensitive to the environmental conditions. A more detailed study comparing mooring line models for a given floating wind turbine design across the full range of operating conditions could provide environmental condition-dependent guidelines for the use of quasi-static versus dynamic mooring line models.

6.4.2 Basis Function Platform Optimization

The work of Chapter 4 took the hydrodynamic basis function technique as far as was reasonable, given certain fundamental limitations in the concept. However, the larger idea of hydrodynamics-based optimization could be worthy of further exploration.

- Development of a more abstract hydrodynamics-based optimization technique could help improve understanding of the design space. Such a technique could use artificial neural networks or meta modelling to delineate constraints on the obtainable hydrodynamic properties and relationships between physical geometry and hydrodynamic coefficients by analyzing a diverse population of platform geometries and calculated hydrodynamic coefficients.

- Development of a geometry-based basis function approach that applies basis function shape modification to the platform geometry, and then analyzes the resulting geometry using WAMIT, could provide a more intuitive alternative to parametric design optimization, while maintaining physical realism.

6.4.3 GA-Based Support Structure Optimization

There is significant potential for further development of the GA-based optimization framework discussed in Chapter 5. Details of some of the specific items for improvement was given in Section 5.7. The greatest potential improvements to the framework are concerned with:

- expanding the design parameterization scheme,
- improving the design modelling and evaluation techniques, and
- increasing the coupling between the design scheme and the modelling techniques through the implementation of an iterative design approach.

These changes involve increased amounts of detail in models of structural design constraints, wind and wave forces, system dynamics, structural loads, manufacturing costs, etc. Accordingly, these changes would be hard to evaluate without increased industrial knowledge to inform the increasingly-detailed models. Further improvements to the framework would therefore be best made in the context of a full-scale floating wind turbine design project, with access to detailed structural and cost data as well as site-specific data from a candidate site under evaluation.

The framework identified consistent trends in the designs found on the Pareto front. Studies of similar four- and seven-cylinder designs are not available in the literature. Using the design trends found by the framework, these support structure configurations could be explored in more detailed design work, leading to new support structure designs that have the potential to be comparable or superior to existing designs.

Bibliography

[1] A. Evans, V. Strezov, and T. J. Evans, "Assessment of sustainability indicators for renewable energy technologies," *Renewable and Sustainable Energy Reviews*, vol. 13, pp. 1082–1088, June 2009.

[2] "World market report," *Renewable Energy World Magazine*, Aug. 2012.

[3] W. Musial and S. Butterfield, "Future for offshore wind energy in the united states," Tech. Rep. 36313, National Renewable Energy Laboratory, Golden, Colorado, June 2004.

[4] J. M. Jonkman, "Dynamics modeling and loads analysis of an offshore floating wind turbine," Tech. Rep. 41958, National Renewable Energy Laboratory, Golden, Colorado, Nov. 2007.

[5] M. Schwartz, D. Heimiller, S. Haymes, and W. Musial, "Assessment of offshore wind energy resources for the united states," Tech. Rep. 45889, National Renewable Energy Laboratory, Golden, Colorado, June 2010.

[6] D. Roddier, C. Cermelli, A. Aubault, and A. Weinstein, "WindFloat: a floating foundation for offshore wind turbines," *Journal of Renewable and Sustainable Energy*, vol. 2, 2010.

[7] B. Skaare, T. D. Hanson, F. G. Nielsen, R. Yttervik, A. M. Hansen, K. Thomsen, and T. J. Larsen, "Integrated dynamic analysis of FLoating offshore wind turbines," tech. rep.

[8] M. A. Lackner and M. A. Rotea, "Passive structural control of offshore wind turbines," *Wind Energy*, 2010.

[9] W. E. Heronemus, "Pollution-Free energy from the offshore winds," in *Preprints, 8th Annual Conference and Exposition, Marine Technology Society, Sept. 11-13*, (Washington, D.C.), 1972.

[10] "Blue h - floating platform technology for offshore wind energy." http://www.bluehgroup.com/. Accessed: May 2011.

[11] H. Stiesdal, "Hywind: The world's first floating MW-scale wind turbine," Dec. 2009.

[12] B. S. Kallesoe, U. S. Paulsen, A. Kohler, and H. F. Hansen, "Aero-Hydro-Elastic response of a floating platform supporting several wind turbines," in *49th AIAA Aerospace Sciences Meeting*, (Orlando, Florida), Jan. 2011.

[13] "Sway | changing the future of wind power." http://www.sway.no/. Accessed: May 2011.

[14] C. Cermelli, A. Aubault, D. Roddier, and T. McCoy, "Qualification of a Semi-Submersible floating foundation for Multi-Megawatt wind turbines," (Houston, Texas), May 2010.

[15] W. J. M. Vijfhuizen, *J,"Design of a Wind and Wave Power Barge,"*. MS dissertation, MS Dissertation, Department of Naval Architecture and Mechanical Engineering, Universities of Glasgow and Strathclyde, Glasgow, Scotland, 2006.

[16] P. Sclavounos, C. Tracy, and S. Lee, "Floating offshore wind turbines: Responses in a seastate pareto optimal designs and economic assessment," Oct. 2007.

[17] L. Vita, U. S. Paulsen, T. F. Pedersen, H. A. Madesn, and F. Rasmussen, "A novel floating offshore wind turbine concept," 2009.

[18] L. Vita, U. S. Paulsen, and T. F. Pedersen, "A novel floating offshore wind turbine concept: New developments," 2010.

[19] C. Cermelli and D. Roddier, "WindFloat: a floating foundation for offshore wind turbines part II: hydrodynamics analysis," in *Proceedings of the ASME 28th International Conference on Ocean, Offshore and Arctic Engineering*, 2009.

[20] G. M. Stewart and M. A. Lackner, "The effect of actuator dynamics on active structural control of offshore wind turbines," *Engineering Structures*, vol. 33, pp. 1807–1816, May 2011.

[21] H. Namik and K. Stol, "Individual blade pitch control of floating offshore wind turbines," *Wind Energy*, vol. 13, pp. 74–85, Jan. 2010.

[22] A. Robertson and J. Jonkman, "Loads analysis of several offshore floating wind turbine concepts," in *Proceedings of the Twenty-first International Offshore and Polar Engineering Conference*, (Maui, Hawaii, USA), 2011.

[23] J. Jonkman and W. Musial, "Offshore code comparison collaboration (OC3) for IEA task 23 offshore wind technology and deployment," Tech. Rep. 48191, National Renewable Energy Laboratory, Golden, Colorado, Dec. 2010.

[24] J. M. Jonkman, S. Butterfield, W. Musial, and G. Scott, "Definition of a 5-MW reference wind turbine for offshore system development," Tech. Rep. 38060, National Renewable Energy Laboratory, Golden, Colorado, Feb. 2009.

[25] D. Roddier, A. Peiffer, J. Weinstein, and A. Aubault, "A generic 5MW WindFloat for numerical tool validation & comparison against a generic 5MW spar," in *Proceedings of the ASME 2011 30th International Conference on Ocean, Offshore and Arctic Engineering*, (Rotterdam, The Netherlands), June 2011.

[26] P. J. Moriarty and A. C. Hansen, "AeroDyn theory manual," tech. rep., NREL, Dec. 2004.

[27] T. Sebastian and M. A. Lackner, "A comparison of First-Order aerodynamic analysis methods for floating wind turbines," (Orlando, Florida), 2010.

[28] J. G. Leishman, "Challenges in modelling the unsteady aerodynamics of wind turbines," *Wind Energy*, vol. 5, no. 2-3, p. 85–132, 2002.

[29] A. Cordle, "State of the art design tools for floating offshore wind turbines," tech. rep., Project UpWind, Mar. 2010.

[30] M. L. Buhl, A. D. Wright, and K. G. Pierce, "FAST_AD code verification: A comparison to ADAMS," (Reno, Nevada), Jan. 2001.

[31] G. Clauss, E. Lehmann, and C. Ostergaard, *Offshore Structures Volume 1: Coneptual Design and Hydromechanics*. Berlin: Springer-Verlag, 1992.

[32] G. I. Taylor, "The spectrum of turbulence," *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, vol. 164, no. 919, p. 476–490, 1938.

[33] M. H. Kim, "Hydrodynamics of offshore structures," in *Developments in Offshore Engineering: Wave Phenomena and Offshore Topics* (J. B. Herbich, ed.), pp. 336–381, Houston, Texas: Gulf Publishing Company, 1999.

[34] J. Falnes, *Ocean Waves and Oscillating Systems*. Cambridge, UK: Cambridge University Press, 2002.

[35] C. Lee, "Wamit thoery manual," tech. rep., MIT Department of Ocean Engineering, Oct. 1995.

[36] A. A. E. Price, C. J. Dent, and A. R. Wallace, "Frequency domain techniques for numerical and experimental modelling of wave energy converters," in *Proceedings of the 8th European Wave and Tidal Energy Conference*, (Uppsala, Sweden), 2009.

[37] G. F. Clauss and L. Birk, "Hydrodynamic shape optimization of large offshore structures," *Applied Ocean Research*, vol. 18, pp. 157–171, Aug. 1996.

[38] E. Wayman, *Coupled Dynamics and Economic Analysis of Floating Wind Turbine Systems*. PhD thesis, MIT, May 2006.

[39] A. Naess and G. M. Ness, "Second-order, sum-frequency response statistics of tethered platforms in random waves," *Applied Ocean Research*, vol. 14, no. 1, pp. 23–32, 1992.

[40] C. W. Hirt and T. B. Nichols, "Volume of fluid (VOF) method for the dynamics of free boundaries," *Journal of Computational Physics*, vol. 39, pp. 201–225, 1981.

[41] R. Dalrymple and B. Rogers, "Numerical modeling of water waves with the SPH method," *Coastal Engineering*, vol. 53, pp. 141–147, Feb. 2006.

[42] M. Masciola, A. Robertson, J. Jonkman, and F. Driscoll, "Investigation of a FAST-OrcaFlex coupling module for integrating turbine and mooring dynamics of offshore floating wind turbines," in *11th International Conference on Offshore Wind Energy and Ocean Energy*, (Beijing, China), Nov. 2011.

[43] M. B. Waris and T. Ishihara, "Influence of mooring force estimation on dynamic response of floating offshore wind turbine system," in *Renewable Energy 2010 International Conference*, (Yokohama, Japan), 2010.

[44] D. Matha, O. Bischoff, U. Fechter, and M. Kuhn, "Non-Linear Multi-Body mooring system model for floating offshore wind turbines," in *EWEA Offshore 2011*, (Amsterdam, The Netherlands), 2011.

[45] J. Azcona, T. A. Nygaard, X. Munduate, and D. Merino, "Development of a code for dynamic simulation of mooring lines in contact with seabed," in *EWEA Offshore 2011*, (Amsterdam, The Netherlands), 2011.

[46] B. J. Buckham and M. J. Nahon, "Formulation and validation of a lumped mass model for low-tension ROV tethers," *International Journal of Offshore and Polar Engineering*, vol. 11, no. 4, 2001.

[47] B. J. Buckham, F. R. Driscoll, M. Nahon, and B. Radanovic, "Torsional mechanics in dynamics simulation of low-tension marine tethers," *International Journal of Offshore and Polar Engineering*, vol. 14, no. 3, pp. 218–226, 2003.

[48] M. Berzeri and A. A. Shabana, "Development of simple models for the elastic forces in the absolute nodal Co-Ordinate formulation," *Journal of Sound and Vibration*, vol. 235, pp. 539–565, Aug. 2000.

[49] R. Driscoll and M. Nahon, "Mathematical modeling and simulation of a moored buoy system," in *OCEANS '96. MTS/IEEE Conference Proceedings*, vol. 1, (Fort Lauderdale, Florida), pp. 517 –523 vol.1, Sept. 1996.

[50] J. M. Jonkman and D. Matha, "Dynamics of offshore floating wind turbines—analysis of three concepts," *Wind Energy*, vol. 14, pp. 557–569, May 2011.

[51] C. C. Tracy, *Parametric design of floating wind turbines.* M.Sc. thesis, MIT, 2007.

[52] D. Matha, *Model Development and Loads Analysis of an Offshore Wind Turbine on a Tension Leg Platform, with a Comparison to Other Floating Turbine Concepts.* PhD thesis, University of Colorada - Boulder, Apr. 2009.

[53] M. Raoof and T. J. Davies, "Determination of the bending stiffness for a spiral strand," *The Journal of Strain Analysis for Engineering Design*, vol. 39, no. 1, pp. 1–13, 2004.

[54] B. S. Kallesoe and A. M. Hansen, "Dynamic mooring line modeling in Hydro-Aero-Servo-Elastic wind turbine simulations," in *Proceedings of the Twenty-first International Offshore and Polar Engineering Conference*, (Maui, Hawaii, USA), June 2011.

[55] J. M. Jonkman, "Definition of the floating system for phase IV of OC3," Technical Report 47535, National Renewable Energy Laboratory, Golden, Colorado, May 2010.

[56] I. Fylling and P. A. Berthelsen, "WINDOPT: an optimization tool for floating support structures for deep water wind turbines," in *Proceedings of the ASME 2011 30th International Conference on Ocean, Offshore and Arctic Engineering*, pp. 767–776, Jan. 2011.

[57] A. R. Henderson, R. Leutz, and T. Fujii, "Potential for floating offshore wind energy in japanese waters," in *Proceedings of The Twelfth (2002) International Offshore and Polar Engineering Conference*, p. 26–31, 2002.

[58] Bulder, van Hees, Henderson, Huijsmans, Pierik, Snijders, Wignants, and Wolf, "Studie naar haalbaarheid van en randvoorwaarden voor drijvende offshore windturbines," tech. rep., TNO, Dec. 2002.

[59] "Vryhof anchor manual," tech. rep., Vryhof Anchors BV, 2012.

[60] R. Ruinen, "Use of drag embedment anchor for floating wind turbines," in *Deep Water Wind Energy Research and Development Planning*, (Washington, D.C.), NREL, 2004.

[61] J. L. Colliat, "Anchors for deepwater to ultradeepwater moorings," vol. OTC14306, (Houston, Texas), 2002.

[62] A. Eltaher, Y. Rajapaksa, and K. Chang, "Industry trends for design of anchoring systems for deepwater offshore structures," vol. OTC 15265, (Houston, Texas), 2003.

[63] "Bruce dennla mk 4 anchor." http://www.bruceanchor.co.uk/Dennla_new_Mk4.html, Aug. 2011.

[64] L. Tao and D. Dray, "Hydrodynamic performance of solid and porous heave plates," *Ocean Engineering*, vol. 35, pp. 1006–1014, July 2008.

[65] L. E. Borgman, "Random hydrodynamic forces on objects," *The Annals of Mathematical Statistics*, vol. 38, no. 1, p. 37–51, 1967.

[66] F. Savenije and J. Peeringa, "Aero-elastic simulation of offshore wind turbines in the frequency domain," tech. rep., Energy Research Center of the Netherlands, 2009.

[67] L. Savenije, *Modeling the dynamics of a spar-type floating offshore wind turbine*. PhD thesis, Technical University Delft, 2009.

[68] M. Hall, "A cumulative Multi-Niching genetic algorithm for multimodal function optimization," *International Journal of Advanced Research in Artificial Intelligence*, vol. 1, Dec. 2012.

[69] I. Y. Kim and O. L. De Weck, "Adaptive weighted-sum method for bi-objective optimization: Pareto front generation," *Structural and multidisciplinary optimization*, vol. 29, no. 2, p. 149–158, 2005.

AppendixAppendix

Appendix A

A Cumulative Multi-Niching Genetic Algorithm for Multimodal Function Optimization

A Cumulative Multi-Niching Genetic Algorithm for Multimodal Function Optimization

Matthew Hall
Department of Mechanical Engineering
University of Victoria
Victoria, Canada

***Abstract*—This paper presents a cumulative multi-niching genetic algorithm (CMN GA), designed to expedite optimization problems that have computationally-expensive multimodal objective functions. By never discarding individuals from the population, the CMN GA makes use of the information from every objective function evaluation as it explores the design space. A fitness-related population density control over the design space reduces unnecessary objective function evaluations. The algorithm's novel arrangement of genetic operations provides fast and robust convergence to multiple local optima. Benchmark tests alongside three other multi-niching algorithms show that the CMN GA has a greater convergence ability and provides an order-of-magnitude reduction in the number of objective function evaluations required to achieve a given level of convergence.**

Keywords- genetic algorithm; cumulative; memory; multi-niching; multi-modal; optimization; metaheuristic.

I. INTRODUCTION

Genetic algorithms provide a powerful conceptual framework for creating customized optimization tools able to navigate complex discontinuous design spaces that could confound other optimization techniques. In this paper, I present a new genetic algorithm that uniquely combines two key capabilities: high efficiency in the number of objective function evaluations needed to achieve convergence, and robustness in optimizing over multi-modal objective functions. I created the algorithm with these capabilities to meet the needs of a very specific optimization problem: the design of floating platforms for offshore wind turbines. However, the algorithm's features make it potentially valuable for any application that features a computationally-expensive objective function and multiple local optima in a discontinuous design space.

Many design optimization problems have computationally-expensive objective functions. While genetic algorithms (GAs) may be ideal optimizers in many ways, a conventional GA's disposal of previously-evaluated individuals from past generations constitutes an unnecessary loss of information. Rather than being discarded, these individuals could instead be retained and used to both inform the algorithm about good and bad regions of the design space and prevent the redundant evaluation of nearly-identical individuals. This could accelerate the optimization process by significantly reducing the number of objective function evaluations required for convergence to an optimal solution.

Examples in the literature of GA approaches that store previously-evaluated individuals in memory to reduce unnecessary or redundant objective function evaluations are sparse. Xiong and Schneider [1] developed what they refer to as a Cumulative GA, which retains all individuals with a high fitness value to use along with the current generation in reproduction. This approach is useful in retaining information about the best regions of the design space, but it does nothing to avoid redundant objective function evaluations. A GA developed by Gantovnik et al. [2], however, does. Their GA stores information about all previous individuals and uses it to construct a Shepard's method response surface approximation of surrounding fitness values, which can be used instead of evaluating the objective function for nearby individuals.

Retaining past individuals to both provide information about the design space and avoid redundant objective function evaluations was my first goal in developing a new GA. My second goal was for the algorithm to be able to identify and converge around multiple local optima in an equitable way.

Identifying multiple local optima is necessary for many practical optimization problems that have multimodal objective functions. Even though an objective function may have only one global optimum, another local optimum may in fact be the preferred choice once additional factors are considered – factors that may be too complex, qualitative, or subjective to be included in the objective function. In the optimization of floating offshore wind turbine platforms, for example, a number of distinct locally-optimal designs exist, ranging from wide barges to deep slender spar-buoys. Though a spar-buoy may have the greatest stability (a common objective function choice), a barge design may be the better choice once ease of installation is considered.

Furthermore, global optimizations often use significant modelling approximations in the objective function for the sake of speed in exploring large design spaces. It is possible for such approximations to skew the design space such that the wrong local optimum is the global optimum in the approximated objective function. In those cases, local gradient-based optimizations with higher-fidelity models in the objective function are advisable as a second optimization stage to verify the locations of the local optima and determine which one of them is in fact the global optimum.

A conventional GA will only converge stably to one local optimum but a number of approaches have been developed for enabling convergence to multiple local optima, a capability referred to as "multi-niching". The Sharing approach, proposed by Holland [3] and expanded by Goldberg and Richardson [4], reduces the fitness of each individual based on the number of neighbouring individuals. The fitness reduction is determined by a sharing function, which includes a threshold distance that determines what level of similarity constitutes a neighbouring individual. A weakness of this approach is that choosing a good sharing function requires a-priori knowledge of the objective function characteristics. As well, the approach has difficulty in forming stable sub-populations, though improvements have been made in this area [5].

An alternative is the Crowding approach of De Jong [6], which features a replacement step that determines which individuals will make up the next generation: for each offspring, a random subset of the existing population is selected and from it the individual most similar to the offspring is replaced by it. Mahfoud's improvement, called Deterministic Crowding [7], removes the selection pressure in reproduction by using random rather than fitness-proportionate selection, and modifies the replacement step such that each crossover offspring competes against the more similar of its parents to decide which of the two enters the next generation.

The Multi-Niche Crowding approach of Cedeño [8] differs from the previous crowding approaches by implementing the crowding concept in the selection stage. For each crossover pair, one parent is selected randomly or sequentially and the other parent is selected as the most similar individual out of a group of randomly selected individuals. This promotes mating between nearby individuals, providing stability for multi-niching. The replacement operation is described as "worst among most similar"; a number of groups are created randomly from the population, the individual from each group most similar to the offspring in question is selected, and the least fit of these "most similar" individuals is replaced by the offspring.

Though the Multi-Niche Crowding approach is quite effective at finding multiple local optima, it and the other approaches described above still provide preferential treatment to optima with greater fitness values. Lee, Cho, and Jung provide another approach, called Restricted Competition Selection [9], that outperforms the previously-mentioned techniques in finding and retaining even weak local optima. In their otherwise-conventional approach, each pair of individuals that are within a "niche radius" of each other are compared and the less fit individual's fitness is set to zero. This in effect leaves only the locally-optimal individuals to reproduce. A set of the fittest of these individuals is retained in the next generation as elites.

Some more recent GAs add the use of directional information to provide greater control of the design space exploration. Hu et al. go so far as to numerically calculate the gradient of the objective function at each individual in order to use a steepest descent method to choose offspring [10]. This approach is powerful, but its large number of function evaluations makes it impractical for computationally-expensive objective functions. Liang and Leung [11] use a more restrained approach in which two potential offspring are created along a line connecting two existing individuals and the four resulting fitness values are compared in order to predict the locations of nearby peaks. By using this information to inform specially-constructed crossover and mutation operators, this algorithm uses significantly fewer function evaluations than other comparable GAs [11].

An approach shown to use even fewer function evaluations is an evolutionary algorithm (EA) by Cuevas and Gonźalez that mimics collective animal behaviour [12]. This algorithm models the way animals are attracted to or repelled from dominant individuals, and retains in memory a set of the fittest individuals. Competition between individuals that are within a threshold distance is also included. Notwithstanding the lack of a crossover function, this algorithm is quite similar in operation to many of the abovementioned GAs and is therefore easily compared with them. It is noteworthy because of its demonstrated efficiency in terms of number of objective function evaluations.

None of the abovementioned multi-niching algorithms retains information about all the previously-evaluated individuals; a GA that combines this sort of memory with multi-niching is a novel creation. In developing such an algorithm, which I refer to as the Cumulative Multi-Niching (CMN) GA, I drew ideas and inspiration from many of the abovementioned approaches. In some cases, I replicated specific techniques, but in different stages of the GA process. The combination of genetic operations to make up a functioning GA is entirely unique.

II. Algorithm Description

The most distinctive feature of the CMN GA is that it is cumulative. Each successive generation adds to the overall population. With the goal of minimizing function evaluations, evaluated individuals are never discarded; even unfit individuals are valuable in telling the algorithm where not to go. The key to making the cumulative approach work is the use of an adaptive proximity constraint that prevents offspring that are overly similar to existing individuals from being added to the population. By using a distance threshold that is inversely proportional to the fitness of nearby individuals, the CMN GA encourages convergence around promising regions of the design space and allows only a sparse population density in less-fit regions of the design space.

This fundamental difference from other GAs enables a number of unique features in the genetic operations of the algorithm that together combine (as summarized in Fig. 2) to make the cumulative multi-niching approach work. The selection and crossover operations are designed to support stable sub-populations around local optima and drive the algorithm's convergence. The mutation operation is designed to encourage diversity and exploration of the design space. The "addition" operation, which takes the place of the replacement operation of a conventional GA, is designed to make use of the accumulated population of individuals in order to avoid redundant or unnecessary fitness function evaluation and guide the GA to produce offspring in the most promising regions of the design space. The fitness scaling operation makes the GA treat local optima equally despite potential

differences in fitness. The details of these operations are as follows.

A. Selection and Crossover

The selection and pairing process for crossover combines fitness-proportionate selection with a crowding-inspired pairing scheme that is biased toward nearby individuals. Whereas Cedeño's Multi-Niche Crowding approach selects the first parent randomly and selects its mate as the nearest of a randomly-selected group, the CMN GA combines factors of both fitness and proximity in its selection operation.

The first parent, P1, of each pair is selected from the population using standard fitness-proportionate selection (FPS) – with the probability of selection proportional to fitness. Then, for each P1, a crowd of N_{crowd} candidate mates is selected using what could be called proximity-proportionate selection (PPS) - with the probability of selection determined by a proximity function describing how close each potential candidate mate, P2, is to P1 in the design space. The most basic proximity function is the inverse of the Euclidean distance:

$$P_{P1,P2} = \frac{1}{\sqrt{\sum_{i=1}^{n}\left(X_i^{P1} - X_i^{P2}\right)^2}} \quad (1)$$

where X is an individual's decision variable vector, with length n. The fittest of the crowd of candidate mates is then selected to pair with P1. This process is repeated for each individual selected to be a P1 parent for crossover.

By having an individual mate with the fittest of a crowd of individuals that are mostly neighbours, mating between members of the same niche is encouraged, though the probability-based selection of the crowd allows occasional mating with distant individuals, providing the important possibility of crossover between niches. This approach contributes to the CMN GA's multi-niching stability and is the basis for crossover-driven convergence of the population to local optima.

In the crossover operation, an offspring's decision variable values are selected at uniform random from the hypercube bounded by the decision variable values of the two parents.

B. Mutation

The mutation operation occurs in parallel with the crossover operation. Mutation selection is done at random, and the mutation of the decision variables of each individual is based on a normal distribution about the original values with a tuneable standard deviation. This gives the algorithm the capability to widely explore the design space. Though individual fitness is not explicitly used in the mutation operation, the addition operation that follows makes it more likely that mutations will happen in fitter regions of the design space.

C. Addition

The cumulative nature of the CMN GA precludes the use of a replacement operation. Instead, an addition operation adds offspring to the ever-expanding population. A proximity constraint ensures that the algorithm converges toward fitter individuals and away from less fit individuals. This filtering, which takes place before the offspring's fitnesses are evaluated, is crucial to the success of the cumulative population approach. By rejecting offspring that are overly similar to existing members of the population, redundant objective function evaluations are avoided.

The proximity constraint's distance threshold, R_{min}, is inversely related to the fitness of the nearest existing individual, $F_{nearest}$, as determined by a distance threshold function. A simple example is:

$$R_{min} = 0.1\,(1.01 - F_{nearest}) \quad (2)$$

This function results in a distance threshold of 0.001 around the most fit individual and 0.101 around the least fit individual, where distance is normalized by the bounds of the design space and fitness is scaled to the range [0 1].

This approach for the addition function allows new offspring to be quite close to existing fit individuals but enforces a larger minimum distance around less fit individuals. As such, the population density is kept high in good regions and low in poor regions of the design space, as determined by the accumulated objective function evaluations over the course of the GA run. A population density map is essentially prescribed over the design space as the algorithm progresses. If the design space was known a priori, the use of a grid-type exploration of the design space could be more efficient, but without that knowledge, this more adaptive approach is more practical.

To adjust for the changing objectives of the algorithm as the optimization progresses – initially to explore the design space and later to narrow in on local optima - the distance threshold function can be made to change with the number of individuals or generation number, G. This can help prevent premature convergence, ensuring all local optima are identified. The distance threshold function that I used to generate the results in this paper is:

$$R_{min} = 0.08\left[1.001 - F_{nearest}(1 - 0.5(0.9)^{G})\right] \quad (3)$$

D. Fitness Scaling

The algorithm described thus far could potentially converge to only the fittest local optimum and not adequately explore other local optima. The final component, developed to resolve this problem and provide equitable treatment of all significant local optima, is a proximity-weighted fitness scaling operation. In most GAs, a scaling function is applied to the population's fitness values to scale them to within normalized bounds and also sometimes to adjust the fitness distribution. A basic approach is to linearly scale the fitness values, F, to the range [0, 1] so that the least fit individual gets a scaled fitness of $F'=0$ and the fittest individual gets a scaled fitness of $F'=1$:

$$F'_i = \frac{F_i - \min(F)}{\max(F) - \min(F)} \quad (4)$$

A scaling function can also be used to adjust the distribution of fitness across the range of fitness values in order to, for example, provide more or less emphasis on moderately-fit individuals. This scaling can be adaptive to the characteristics of the population. For the results presented

here, I used a second, exponential scaling function to adjust the scaled fitness values so that the median value is 0.5:

$$F''_i = (F'_i)^{\left[\frac{\ln(0.5)}{\ln(\text{median}(F'))}\right]} \quad (5)$$

Proximity-weighted fitness scaling, a key component of the CMN GA, adds an additional scaling operation. This operation relies on the detection of locally-optimal individuals in the population. The criterion I used, for simplicity, is that an individual is considered to represent a local optimum if it is fitter than all of its nearest N_{min} neighbours. In the proximity-weighted fitness scaling operation, scaling functions (4) and (5) are applied multiple times to the population, each time normalizing the results to the fitness of a different local optimum. So if m local optima have been identified, each individual in the population will have m scaled fitness values. These scaled fitness values F'' are then combined for each individual i according to the individual's proximity to each respective local optimum j to obtain the population's final scaled fitness values:

$$F'''_i = \frac{\sum_{j=1}^{m} P_{i,j} F''_{i,j}}{\sum_{j=1}^{m} P_{i,j}} \quad (6)$$

Proximity, $P_{i,j}$, can be calculated as in (1). This process gives each local optimum an equal scaled fitness value, as is illustrated for a one-dimensional objective function in Fig. 1.

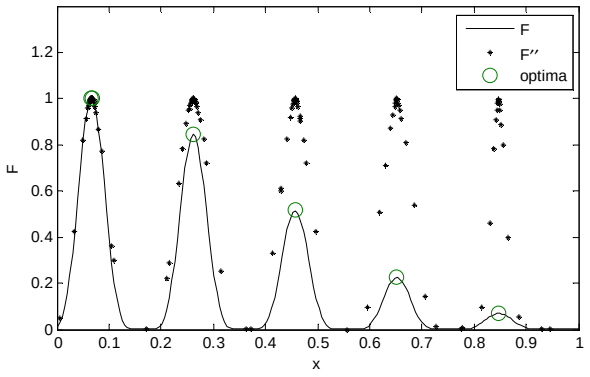

Figure 1. Proximity-weighted fitness scaling.

E. CMN GA Summary

Fig. 2 describes the overall structure of the CMN GA, outlining how the algorithm's operations are ordered and how the addition operation filters out uninformative offspring. The next section demonstrates the algorithm's effectiveness at multi-niche convergence with a minimal number of objective function evaluations.

III. Performance Results

To benchmark the CMN GA's performance, I tested it alongside three other multi-niching algorithms on four generic multimodal objective functions. These four multimodal functions have been used by many of the original developers of multi-niching GAs [8].

```
Step 0: (Initialization)
  Randomly generate N_pop individuals
  Evaluate the individuals' fitnesses F

Step 1: (Fitness Scaling)
  Calculate distances between individuals
  Identify locally-optimal individuals
  For each individual i:
    For each locally-optimal individual j:
        Calculate scaled fitness F''_i,j
    Calculate proximity-weighted fitness F'''_i

Step 2: (Crossover)
  Select a P1 from the population using FPS
  Select a crowd of size N_crowd using PPS
  Select the fittest in the crowd to be P2
  Cross P1 and P2 to produce an offspring
  If offspring satisfies distance threshold:
    Add to population and calculate fitness F
  Repeat N_try times or until N_crossover offspring
  have been added to the population

Step 3: (Mutation)
  Randomly select a mutation individual
  Mutate individual to produce an offspring
  If offspring satisfies distance threshold:
    Add to population and calculate fitness F
  Repeat N_try times or until N_mutate offspring
  have been added to the population

Step 4: (New Generation)
  Repeat from Step 1 until stopping criterion
  is met
```

Figure 2. CMN GA outline.

The first, F1, is a one-dimensional function featuring five equal peaks, shown in Fig. 3.

$$F_1(x) = \sin^6(5.1\,\pi x + 0.5) \quad (7)$$

The second, F2, modifies F1 to have peaks of different heights, shown in Fig. 4.

$$F_2(x) = \exp\left(-\frac{4(ln2)(x-0.0667)^2}{0.64}\right) F_1(x) \quad (8)$$

The third, F3, is a two-dimensional Shekel Foxholes function with 25 peaks of unequal height, spaced 16 units apart in a grid, as shown in Fig. 5.

$$F_3(x,y) = 0.002 + \sum_{i=1}^{25} \frac{1}{i+(x-A_i)^6+(y-B_i)^6} \quad (9)$$

The fourth, F4, is an irregular function with five peaks of different heights and widths, as listed in Table 1 and shown in Fig. 6.

$$F_4(x,y) = \sum_{i=1}^{5} \frac{H_i}{1+W_i[(x-A_i)^2+(y-B_i)^2]} \quad (10)$$

In F3 (9) and F4 (10), A_i and B_i are the x and y coordinates of each peak. In F4 (10), H_i and W_i are the height and width parameters for each peak.

These four functions test the algorithms' multi-niching capabilities in different ways.

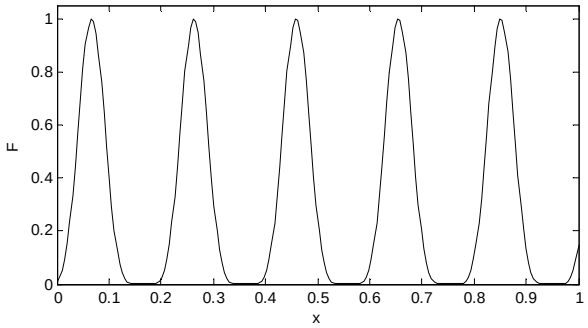

Figure 3. F1 objective function.

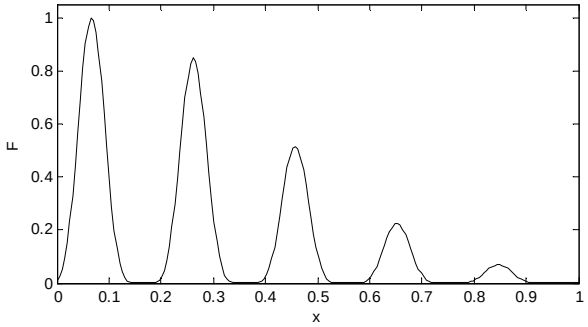

Figure 4. F2 objective function.

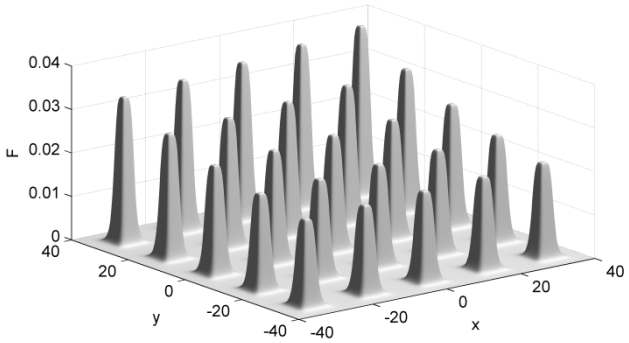

Figure 5. F3 objective function.

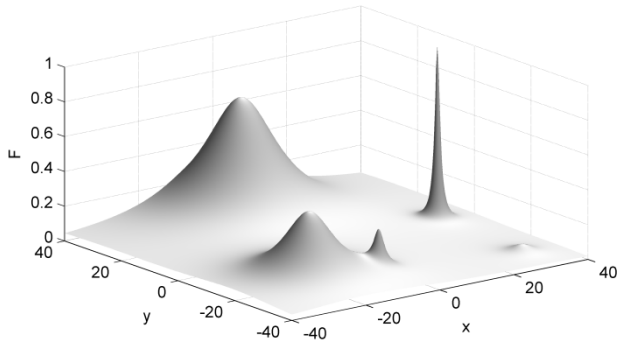

Figure 6. F4 objective function.

TABLE I. F4 OBJECTIVE FUNCTION PEAKS

I	A_i	B_i	H_i	W_i
1	-20	-20	0.4	0.02
2	-5	-25	0.2	0.5
3	0	30	0.7	0.01
4	30	0	1.0	2.0
5	30	-30	0.05	0.1

The two other multi-niching GA approaches I compare the CMN GA against are Multi-Niche Crowding (MNC) [8] and Restricted Competition Selection (RCS) [9]. I chose these two because they are very well-performing examples of two different approaches to GA multi-niching. I implemented these techniques into a GA framework that is otherwise the same as the CMN GA in terms of how it performs the crossover and mutation operations. Crossover offspring decision variable values are chosen at uniform random from the intervals between the decision variables of the two parents. Mutation offspring decision variable are chosen at random using normal distributions about the unmutated values with standard deviations of 40% of the design space dimensions.

For further comparison, I also implemented the Collective Animal Behaviour (CAB) evolutionary algorithm [12]. It is a good comparator because it has many common features with multi-niching GAs, but has been shown to give better performance than many of them, particularly in terms of objective function evaluation requirements.

The values of the key tunable parameters used in each algorithm are given in Tables 2 to 5. N_{pop} describes the population size, or the initial population size in the case of the CMN GA. For the RCS GA, N_{elites} is the number of individuals that are preserved in the next generation. I tuned the parameter values heuristically for best performance on the objective functions. For the MNC, RCS, and CAB algorithms, I began by using the values from [8], [9], and [12], respectively, but found that modification of some parameters gave better results. The meanings of the variables in Table 4 can be found in [12].

To account for the randomness inherent in the operation of a genetic or evolutionary algorithm, I ran each algorithm ten times on each objective function to obtain a reliable characterization of performance. The metric I use to measure the convergence of the algorithms to the local optima is the sum of the distances from each local optimum X^*_j to the nearest individual. By indicating how close the algorithm is to identifying all of the true local optima, this aggregated metric represents what is of greatest interest in multimodal optimization applications. The assumption is that in real applications it will be trivial to determine which evaluated individuals represent local optima without a-priori knowledge of the objective function.

TABLE 2. PARAMETERS FOR THE MNC GA TECHNIQUE

Function	F1 & F2	F3 & F4
N_{pop}	50	200
$N_{crossover}$	45	180
$N_{mutation}$	5	20
C_S	15	75
C_F	3	4
S	15	75

TABLE 3. PARAMETERS FOR THE RCS GA TECHNIQUE

Function	F1 & F2	F3 & F4
N_{pop}	10	80
N_{elites}	5	30
$N_{crossover}$	8	50
$N_{mutation}$	2	30
R_{niche}	0.1	12

TABLE 4. PARAMETERS FOR THE CAB EA TECHNIQUE

Function	F1 & F2	F3 & F4
N_{pop}	20	200
B	10	100
H	0.6	0.6
P	0.8	0.8
v	0.01	0.001
ρ	0.1	4

TABLE 5. PARAMETERS FOR THE CMN GA TECHNIQUE

Function	F1 & F2	F3 & F4
$N_{pop\ (initial)}$	10	100
$N_{crossover}$	3	20
$N_{mutation}$	2	12
N_{min}	3	6
N_{crowd}	10	20
N_{try}	100	100

Figures 7 to 10 show plots of the convergence metric versus the number of objective function evaluations for each optimization run. Using these axes gives an indication of algorithm performance in terms of my two objectives for the CMN GA, convergence to multiple local optima and minimal objective function evaluations. Figures 7, 8, 9, and 10 compare the performance of each algorithm for objective functions F1, F2, F3, and F4, respectively.

In the results for objective function F4, the MNC and CAB algorithms consistently failed to identify the shallowest peak. Accordingly, I excluded this peak from the convergence metric calculations for these algorithms in the data of Fig. 10 in order to provide a more reasonable view of these algorithms' performance. The CMN GA also missed this peak in one of the runs, as can by the one anomalous curve in Fig. 10, wherein the convergence metric stagnates at a value of 2. As is the case with other multi-niching algorithms, missing subtle local optima is a weakness of the CMN GA, but it can be mitigated

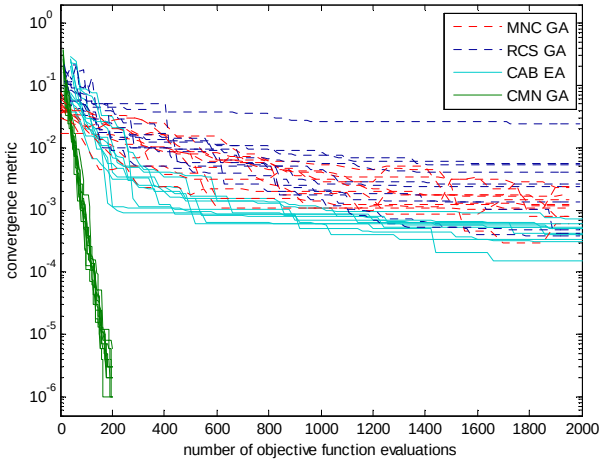

Figure 7. GA performance for F1 objective function runs.

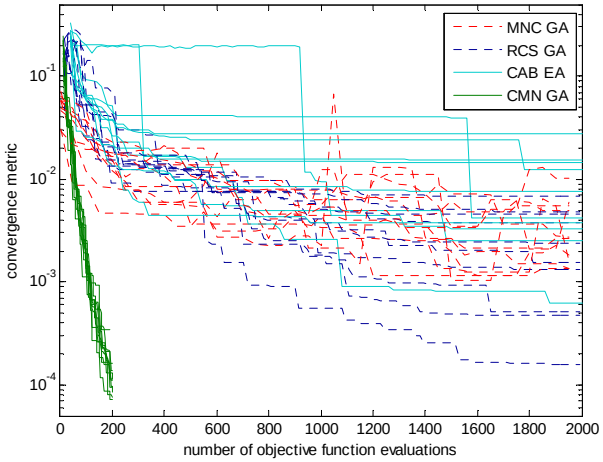

Figure 8. GA performance for F2 objective function runs.

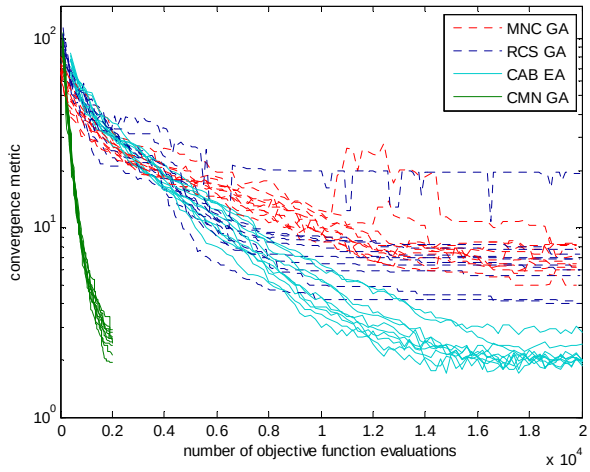

Figure 9. GA performance for F3 objective function runs.

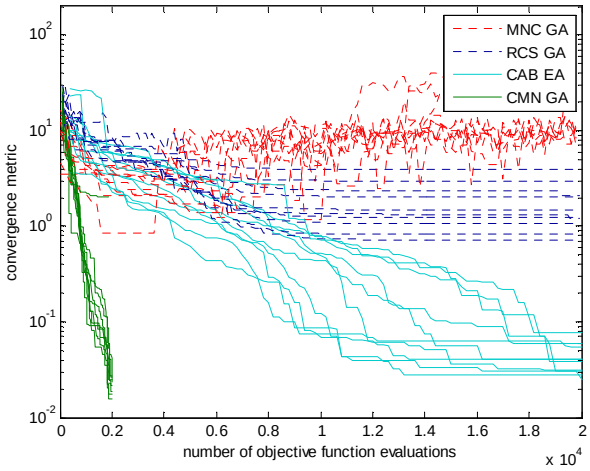

Figure 10. GA performance for F4 objective function runs.

by careful choice of algorithm parameters and verifying results through multiple optimization runs.

Fig. 11 is a snapshot of a population generated by the CMN GA on the F4 objective function. The distribution of the 1000 individuals in the figure illustrates how the algorithm clearly identifies the five local optima and produces a high population density around them regardless of how shallow or sharp they may be. Fig 12 shows how, with the same input parameters, the CMN GA is just as effective with the 25 local optima of the F3 objective function.

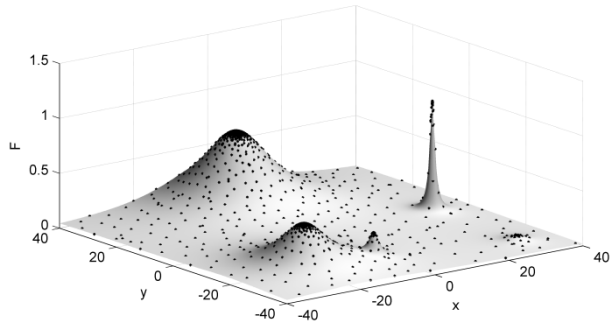

Figure 11. CMN GA exploration of F4 objective function.

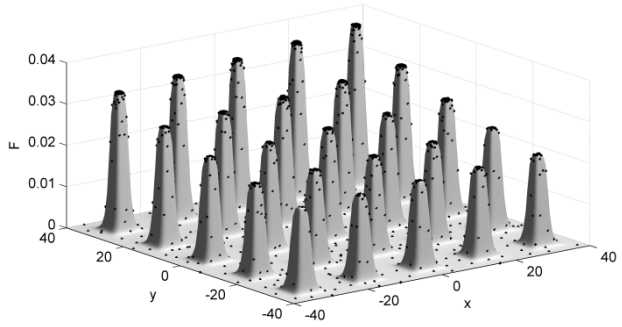

Figure 12. CMN GA exploration of F3 objective function.

Though more rigorous tuning of parameters could result in slight performance improvements in any of the four algorithms I compared, the order-of-magnitude faster convergence of the CMN GA gives strong evidence of its superior performance in terms of multimodal convergence versus number of objective function evaluations.

It should be noted that this measure of performance, reflective of the design goals of the CMN GA, is only indicative of performance on optimization problems where evaluating the objective function dominates the computational effort. The algorithm operations of the CMN GA are themselves much slower than those of the other algorithms, so the CMN GA could be inferior in terms of computation time on problems with easily-computed objective functions. As well, with its ever-growing population, the CMN GA's memory requirements are greater than those of the other algorithms. In a sense, my choice of measure of performance puts the MNC, RCS, and CAB algorithms at a disadvantage because, unlike the CMN GA, these algorithms were not designed specifically for computationally-intensive objective functions. That said, convergence versus number of function evaluations is the most relevant measure of performance for optimizing over computationally-expensive multimodal objective functions, and the algorithms I chose for comparison represent three of the best existing options out of the selection of applicable GA/EA approaches available in the literature.

IV. Conclusion

In the interest of efficiently finding local optima in computationally-expensive objective functions, I created a genetic algorithm that converges robustly to multiple local optima with a comparatively small number of objective function evaluations. It does so using a novel arrangement of genetic operations in which new individuals are continuously added to the population; I therefore call it a Cumulative Multi-Niching Genetic Algorithm. The tests presented in this paper demonstrate that the CMN GA meets its goals – convergence to multiple local optima with minimal objective function evaluations – strikingly better than alternative genetic or evolutionary algorithms available in the literature. It therefore represents a useful new capability for optimization problems that have computationally-expensive multimodal objective functions. The proximity constraint approach used to control the accumulation of individuals in the population may also be applicable to other metaheuristic algorithms.

V. References

[1] Y. Xiong and J. B. Schneider, "Transportation network design using a cumulative genetic algorithm and neural network," *Transportation Research Record*, no. 1364, 1992.

[2] V. B. Gantovnik, C. M. Anderson-Cook, Z. Gürdal, and L. T. Watson, "A genetic algorithm with memory for mixed discrete–continuous design optimization," *Computers & Structures*, vol. 81, no. 20, pp. 2003–2009, Aug. 2003.

[3] J. H. Holland, *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence.* U Michigan Press, 1975.

[4] D. E. Goldberg and J. Richardson, "Genetic algorithms with sharing for multimodal function optimization," in *Proceedings of the Second International Conference on Genetic Algorithms and their Application*, 1987, pp. 41–49.

[5] B. L. Miller and M. J. Shaw, "Genetic algorithms with dynamic niche sharing for multimodal function optimization," in *Proceedings of IEEE International Conference on Evolutionary Computation*, 1996, pp. 786–791.

[6] K. A. De Jong, "Analysis of the behavior of a class of genetic adaptive systems," PhD Thesis, University of Michigan, 1975.

[7] S. W. Mahfoud, "Crowding and preselection revisited," *Parallel problem solving from nature*, vol. 2, pp. 27–36, 1992.

[8] W. Cedeño, "The multi-niche crowding genetic algorithm: analysis and applications," PhD Thesis, University of California Davis, 1995.

[9] C.-G. Lee, D.-H. Cho, and H.-K. Jung, "Niching genetic algorithm with restricted competition selection for multimodal function optimization," *Magnetics, IEEE Transactions on*, vol. 35, no. 3, pp. 1722 –1725, May 1999.

[10] Z. Hu, Z. Yi, L. Chao, and H. Jun, "Study on a novel crowding niche genetic algorithm," in *2011 IEEE 2nd International Conference on Computing, Control and Industrial Engineering (CCIE)*, 2011, vol. 1, pp. 238 –241.

[11] Y. Liang and K.-S. Leung, "Genetic Algorithm with adaptive elitist-population strategies for multimodal function optimization," *Applied Soft Computing*, vol. 11, no. 2, pp. 2017–2034, Mar. 2011.

[12] E. Cuevas and M. González, "An optimization algorithm for multimodal functions inspired by collective animal behavior," *Soft Computing*, Sep. 2012.

Appendix B

Genetic Algorithm Implementation Details and Settings

B.1 Treatment of Discontinuities

In the geometry scheme of Section 5.2, as in many optimization problems dealing with multiple configuration options, some configurations or design points can be described by more than one combination of decision variables. In this design space, the potential for overlap arises from the multiple ways a cylinder could be made to not exist.

The inner cylinder could be eliminated by either its radius or its draft going to zero. The outer cylinders could be eliminated by their radii, drafts, spacing, or number going to zero.

These overlaps or redundancies in the design space have the potential to slow down or confuse both the optimization algorithm and the person interpreting its results. It is therefore useful to try to choose the decision variables so as to make the design space as logical as possible, as was done in the framework. Even then, regions of design space overlap still exist. Techniques can be applied to collapse such overlapping regions of the design space – meaning that (1) points in the overlap region are now only described by one set of decision variables and (2) the optimization algorithm no longer treats an overlap region as multiple regions. This is important, because the GA operates on measurements of distance and counting, so potential double-counting or not-counting of individuals would be a problem.

The technique that was implemented to deal with this relies on the use of NaNs in non-applicable variables to signal to the GA that these dimensions are to be ignored

in distance calculations. Either of the radius variables, R_I or R_O, is set to zero if the respective cylinder(s) is/are disabled. The remaining design variables associated with the disabled cylinder(s) are given a NaN value, which flags several special behaviours in the GA:

- In design space distance calculations, these dimensions are not considered. (So, for example, a single-cylinder design will actually fill the entire hypervolume defined by the disabled outer cylinder variables (N_F, R_F, H_O, and R_{HP}).)
- In crossover and mutation operations, each of the disabled variables is assigned a value equal to it's lower bound, in case the disabled cylinder(s) become enabled in the offspring.

With this strategy, the potential for redundant or overlapping design space regions is eliminated.

The integer-valued number of outer cylinders variable, N_F, is dealt with in the GA by simple rounding, and bounds set at 0.5 above and 0.5 below the desired upper and lower integer bounds on the variable.

B.2 Treatment of Constraints

The majority of the constraints in the design space are nonlinear, and many of the constraints are performance-dependent. To handle this, the GA was designed to handle constraints by simply discarding a non-complying individual and generating a new candidate individual in replacement. For greatest efficiency, compliance with constraints is evaluated as soon as possible in the design evaluation process. This means that some constraints, such as hydrostatic stability and net buoyancy, are checked as soon as the mass and mooring models have been applied, while other constraints, such as maximum dynamic pitch angle and potential for snap loading of mooring lines, are checked only after the full hydrodynamic analysis has been completed.

The procedure used for calculating the performance of a candidate support structure based on its design variables is as follows:

- design variables generated
- design variable constraints checked (things like collisions)

- call to C++ DLL
  - platform surface discretized
  - mass properties calculated
  - float truss connections sized, and system mass updated
  - mooring system analized a first time (maximum loads identified)
  - sufficient buoyancy constraint checked
  - mooring line tendons sized, and system mass updated
  - mooring system analyzed a second time
  - sufficient buoyancy constraint checked
- ballast added (if applicable)
- final mass properties calculated
- static pitch constraint checked
- cost constraint checked
- net buoyancy constraint checked (again)
- call to WAMIT (including checks for errors)
- all linear matrices loaded
- RAOs evaluated across frequency spectrum for each metocean condition (including iterative solution of linearized viscous drag)
- objective function calculated

B.3 GA Settings and Functions

The main parameters for the GA's operation, which are described in detail in Appendix A, are listed in Table B.1 along with the values prescribed to them for the different optimization scenarios. An additional parameter, minimum distance between local optima (d_{opt}), has been added to the algorithm. All distances are normalized to the bounds of the design space.

Table B.1: GA settings

Setting	Single cyl.	Full space	Description
$N_{population}$	20	100	initial population size
$N_{crossover}$	8	8	number of crossover offspring
$N_{mutation}$	20	20	number of mutation offspring
N_{min}	4	9	no. of nearest neighbours for optima detection
N_{crowd}	20	20	selection pairing crowd size
N_{try}	10000	10000	maximum offspring generation attempts
d_{opt}	0.2	0.2	minimum distance between local optima
$F_{scaling}$	see (B.1)	see (B.1)	fitness scaling function
R_{min}	see (B.2)	see (B.3)	distance threshold for new individuals

The fitness scaling function used in the optimizations is

$$F_i'' = (F_i')^{\frac{\log(0.5)}{\log(median(F'))}} \tag{B.1}$$

where F_i'' is the scaled fitness and F_i' is the normalized fitness of individual i.

The same distance threshold function is used for both mutation and crossover offspring. For the single-cylinder design space it is

$$R_{min} = (1 - 0.9(F''_{nearest})^5)0.999^G, \tag{B.2}$$

and for the full design space it is

$$R_{min} = (1 - 0.9(F''_{nearest})^5)0.9995^G, \tag{B.3}$$

where $F''_{nearest}$ is the scaled fitness value of the nearest individual, and G is the generation number.

Appendix C

Comparison of Framework Model Results to Published Data

C.1 Comparison Description

Many of the input settings in the global optimization framework were set to calibrate the results to be consistent with two existing designs: the Hywind and the WindFloat. Details of these two designs are readily available, which makes them good sources of comparison to provide a cursory verification of the frequency-domain model used in the framework. Response amplitude operators for both designs can be found in the OC4 WindFloat specifications paper [25].

Designs were made within the parameterization of the framework to approximate the OC3 Hywind and OC4 WindFloat designs. Some specifications of these designs are compared in Table C.1. Differences in the Hywind comparison can be attributed to the platform parameterization used in the framework. The framework assumes the taper happens from 1/8 to 1/4 of the draft, while the taper on the OC3 Hywind design occurs from 1/30 to 1/10 of the draft; this causes the mass, displacement, and moments of inertia of the framework design to be less than the OC3 Hywind design for the same overall draft and diameters. The much smaller differences in the WindFloat comparison can be attributed to the framework setting for hull thickness, heave plate thickness, and connective element sizing resulting in slightly different values from those used in the OC4 WindFloat design.

In the figures that follow, the response amplitude operators for surge, heave, pitch, and nacelle acceleration are compared between the results generated in the framework

Table C.1: Comparison of support structure properties

	Hywind		WindFloat	
	OC3 specs [55]	Framework	OC4 Specs [25]	Framework
draft (m)	120	120	17	17
$m_{platform}$ (tonnes)	7466	6817	4640	4300
$\forall$ (m^3)	8029	7493	5	4606
CG draft (m)	89.9	90.6	3.73	
I_4 (kg-m^2)	6.46E+10	6.24E+10	5.72E+09	5.62E+09
I_5 (kg-m^2)	6.46E+10	6.24E+10	5.65E+09	5.52E+09
I_6 (kg-m^2)	1.64E+08	9.36E+07	3.26E+06	2.31E+09

and the results provided in [25].

C.2 Discussion

The levels of disagreement between most of the RAO curves are not unreasonable, given the differences in physical properties evidenced in Table C.1 as well as the different treatments of viscous drag forces between the framework and the modelling tools that were used to produce the curves in [25]. The notable exception is Figure C.1(a), which shows extreme disagreement between the framework's calculation of the Hywind surge RAO compared to the published data. While the aforementioned difference in taper between these designs would account for some differences in the hydrodynamics, no explanation has yet been made for the extreme level of disagreement shown in the figure.

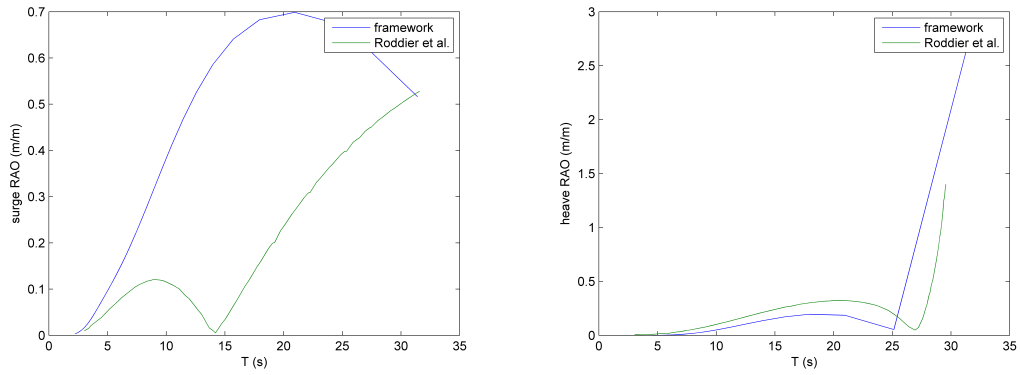

(a) Surge

(b) Heave

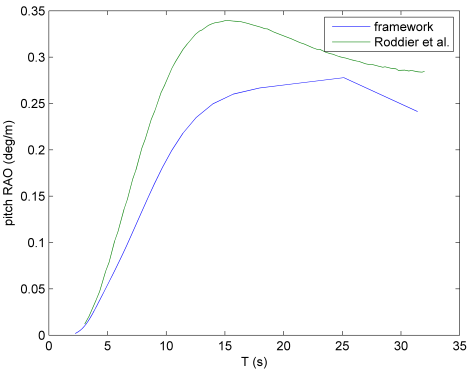

(c) Pitch

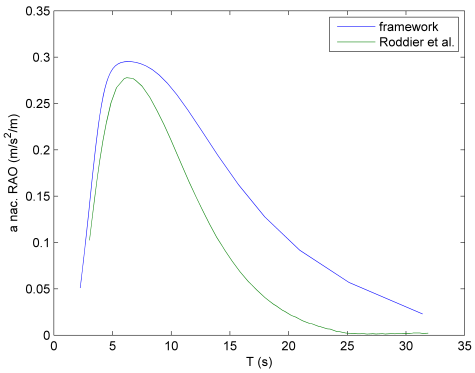

(d) Nacelle Acceleration

Figure C.1: Hywind RAO comparison

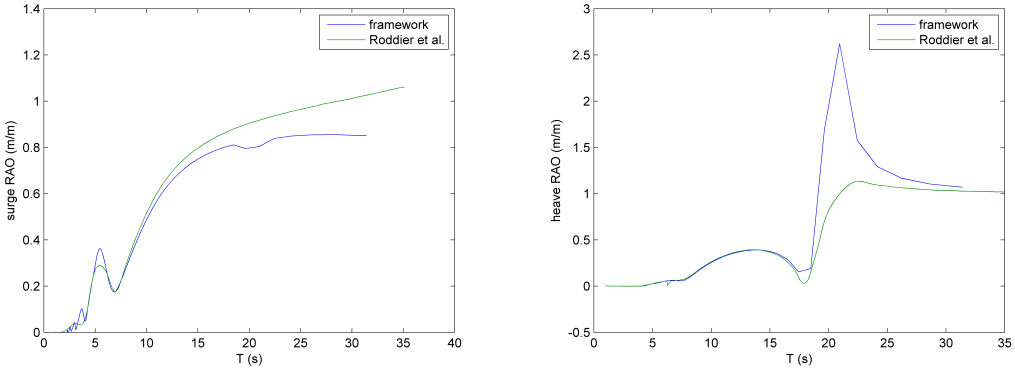

(a) Surge

(b) Heave

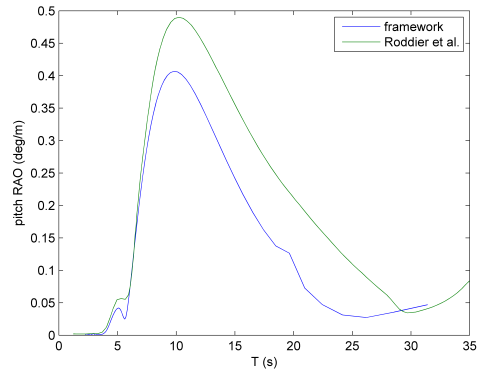

(c) Pitch

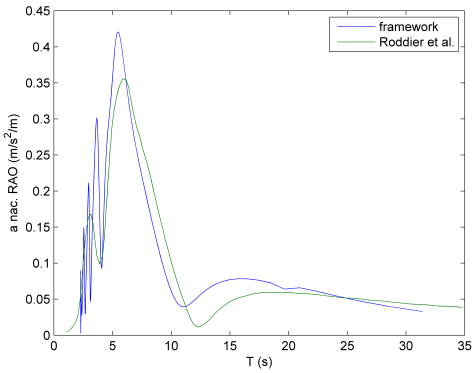

(d) Nacelle Acceleration

Figure C.2: WindFloat RAO comparison