MSc Thesis

Numerical investigation of control methods for reduction of motion of suspended jackets in offshore transport

Jerin Thomas

Student number: 6276423

Delft University of Technology

TU Delft

(*Page intentionally left blank*)

Numerical investigation of control methods for reduction of motion of suspended jackets in offshore transport

by

Jerin Thomas
6276423

In partial fulfilment of the requirements for the degree of
Master of Science in Civil Engineering *at the*
Delft University of Technology

In collaboration with ***Allseas***

An electronic version of this thesis is available at
`http://repository.tudelft.nl/`.

(*Page intentionally left blank*)

Graduation Committee

The graduation committee consists of the following members:

Dr. Ir. P.C. (Peter) Meijers	TU Delft
Dr. Ir. Panagiota Atzampou	TU Delft
Dr. Ir. A.B. Faragau	TU Delft
Dr. Ir. Niels Mallon	Allseas Engineering B.V.
Ir. Jochem de Kwant	Allseas Engineering B.V.

(*Page intentionally left blank*)

Preface

This thesis marks the conclusion of my Master of Science in Civil Engineering at the Delft University of Technology, specializing in Structural Engineering. My research explores the analytical modeling and numerical simulation of offshore heavy-lift operations, specifically comparing hoist geometry optimization and active tugger damping to expand operational weather windows for suspended jacket transportation.

This work was conducted during my graduation internship at Allseas Engineering B.V.. I would like to thank my company supervisors, Dr. Ir. Niels Mallon and Ir. Jochem de Kwant, for providing the opportunity to investigate this topic and for their guidance throughout the project.

I am also deeply grateful to my professors at TU Delft for equipping me with the foundational knowledge in structural mechanics and dynamics needed to proudly call myself an engineer. In particular, I want to express my sincere appreciation to my academic supervisors, Dr. Ir. Panagiota Atzampou and Dr. Ir. Peter Meijers, for their invaluable feedback, continuous direction, and unwavering support.

Finally, none of this would have been possible without the people closest to me. I express my deepest gratitude to my parents for laying the foundation for my ambitions and making this entire journey possible. Thank you as well to my family and friends for standing by me through the ups and downs of my studies.

Jerin Thomas

Delft, August 2026

Abstract

Marine operations are fundamentally constrained by safety, efficiency, and cost, factors that become critically important for massive heavy-lift vessels like the *Pioneering Spirit*. This vessel utilizes a bespoke Jacket Lift System (JLS) engineered for offshore jacket removal and installations, which requires transporting large steel jackets in a freely suspended state. During open-ocean transit, wave excitation drives severe pendular sway in the payload, restricting the workable weather window. Here we develop a simplified, computationally efficient 4-degree-of-freedom model to simulate these transit dynamics and evaluate methods to reduce jacket motion. We show that passive hoist optimisation, achieved by altering the rigging geometry, substantially reduces resonant motion and improves operability without requiring any additional hardware. The performance of this passive detuning is directly compared against an active, velocity-proportional tugger damping system. Overall, this work establishes the reduced-order model as a practical tool for route-specific rigging assessment. It proves that passive geometry reconfiguration is the most effective immediate mitigation, while clearly defining the strict conditions under which high-capacity active tuggers become beneficial.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The offshore heavy-lift industry deals with the installation and removal of some of the largest structures ever built at sea such as fixed platforms and jackets to the foundations of the offshore wind farms reshaping the global energy mix. For these heavy payloads specialised vessels are used increasingly to execute the lifting and transit operations. But the envorimental loads such as waves and winds in oceans affect these operations by reducing the safety and efficiency. Allseas developed the Pioneering Spirit with a bespoke jacket lifting system specifically designed for jacket removal and installation operations. The Jacket Lifting System (JLS) uses four hoists connected to long, skiddable beams at the aft of the vessel. During transit, the jacket can begin to oscillate as the beams transmit vessel motion through their large lever arm. The resulting motion of such massive payloads can be detrimental to marine operations. Understanding and controlling this motion is therefore a central determinant of whether a given offshore heavy-lift operation can proceed at all, and it forms the basis of this thesis. This chapter positions the present work within that context by setting out the research problem (Section 1.2), the research objectives (Section 1.3), the research scope (Section 1.4), the research questions (Section 1.5), the research theory and methods (Section 1.6), and the structure of the report (Section 1.7).

Figure 1.1: Pioneering Spirit

1.1 Research Context

The offshore industry has two large markets that are simultaneously growing: the installation and decommissioning of offshore oil and gas infrastructure, and the rapidly expanding offshore-wind sector. Both rely on purpose-built crane vessels capable of lifting topsides, jackets, and turbine substructures. Allseas, established as a marine contractor, has developed and now operates the *Pioneering Spirit*: the largest construction vessel in the world by displacement. Originally conceived for single-lift platform installation and removal, the vessel has been progressively extended into jacket installation through the Jacket Lift System (JLS). The jacket lift system is located at the aft of the vessel and consists of two 170 meter tilting lift beams, hinged around the stern. During a jacket removal, the jacket is cut-off at the bottom, hoisted from the tip of the beams, partially rotated, and then tilted inboard. Jacket removal operations and transport to barge or ports will take place in open sea and are thus subjected to environmental actions.

Once hoisted, the jacket must be transported to a barge or disposal port in open seas. During this transit phase, the JLS can configure the payload in three ways: fully tilted to horizontal (resting entirely on the beams), partially tilted (resting against the beams at approximately 70°), or freely suspended (with no contact against the beams). The transit phase of a JLS lift exposes the suspended jacket to vessel motion as the jacket is free to move relative to the vessel in the suspended arrangement. The Pioneering Spirit, like all vessels, responds to wave excitation in all six rigid-body degrees of freedom, with peak motion amplitudes occurring near the vessel's own natural periods. The lever arm between the vessel CoG and the JLS beams which holds the hoist wires is large (approximately 215 m longitudinal and 100 m vertical), so even small vessel rotations produce significant translations at the suspension point.These hinge motions drive the suspended jacket as a pendulum.

Large sway amplitudes increase the risk of collision between the jacket and the lift beams, generate load variations in the lift wires, and make the operation more difficult to control. Allseas aims to apply a limit of 3 m on the largest relative jacket sway expected in a three-hour exposure to make sure the marine guidelines for operations are maintained, and this limit is used throughout the present work as the criterion for whether a sea state is workable. These effects restrict the workability envelope—the set of sea states for which safe transportation can be maintained and thereby impact asset throughout. Motion reduction control of the suspended jacket can therefore improve operational robustness and expand the permissible operating conditions.

1.2 Research Problem

For jacket transport on the Pioneering Spirit, the dominant operational risk in the transit phase is dynamic excitation of the suspended jacket by wave-induced vessel motion. The mechanism is well understood qualitatively but is not yet quantitatively characterised in a form suitable for control oriented studies.

Existing in-house analyses of the operation rely on the full nonlinear OrcaFlex model of the vessel and lift system along with the payloads. This model captures the relevant physics with high fidelity but is computationally expensive. A single sea-state simulation of representative duration consumes time on the order of minutes or hours, and parametric studies in which the rigging arrangement, hoist-wire pretensions, or control gains are varied across realistic ranges quickly become infeasible. As a consequence, the design of any active or passive controls to improve workability is currently limited.

A control-oriented *reduced-order model* (ROM) of the system would close this gap. A ROM that captures the dominant dynamics of the suspended jacket in swell waves could be evaluated thousands of times in a parametric study at negligible computational cost, and would help quantify the physical structure of the response (modes, frequencies, transfer functions) in a form applicable to control studies. The full nonlinear model would then be reserved for the final verification or validations.

No such reduced-order model exists for this specific configuration and the Jacket Lifting System although there are studies on offshore cranes which uses such reduced models for studies. The combination of bespoke lifting system on the PS and the kinematic transfer from vessel CoG to the JLS suspension point through a long lever arm produces a coupled multi-degree-of-freedom system whose dynamics cannot be deduced from generic pendulum-on-a-vessel results seen in studies. The first expected contribution of this thesis is to construct and validate this model.

With a validated model in hand, two improvement ways of reducing the jacket motion are then explored. The first — *hoist-arrangement optimisation* — exploits the fact that the natural periods of the system depend on the rigging geometry , of which the operator can vary within physical limits. Changing the rigging geometry moves the sway period further from the wave periods and also changes how much of the suspension-point motion is passed on to the jacket. The second path — *tugger control* — adds tugger lines that apply a force opposing the sway velocity, so removing energy from the motion. The two are compared separately and in combination, so that the contribution of each can be identified. Both avenues require a control-oriented model of the underlying plant. This is the gap the present work addresses.

1.3 Research Objective

The problem statement identifies that improving the motion control of the suspended jacket during JLS transit requires a control-oriented model capable of fast evaluation. Therefore, the overall objective of this research is defined as:

> *Develop and validate a control-oriented reduced-order model for the suspended jacket dynamics during JLS transit on the Pioneering Spirit through open sea waves, and use that model to evaluate motion-reduction strategies that increase the operational workability envelope.*

1.4 Research Scope

The scope of this research is limited to the transit phase of the jacket lift operation which is strictly in air, and aims to develop a numerical model to predict the system's dynamic response with and without active tugger control. The model represents the suspended jacket as a simplified four-degree-of-freedom rigid body (sway, heave, roll, and yaw). More specifically, effects from surge and pitch are neglected due to vessel dynamic positioning constraints and minimal excitation in transit operations, thereby focusing on the dominant pendular dynamics. The model considers first-order swell wave loading as the major disturbance .Secondary environmental loads, such as wind, current, slamming, and second-order wave drift, are small relative to the first-order swell excitation and are therefore not included in the present analysis. Jacket transport operations encounter varying swell conditions that can trigger severe resonant motions. This research will specifically focus on the operational transit limits defined by the operator, and also on ways to expand it using the help of control methods. The analysis conservatively assumes pure beam seas (90°) to evaluate the worst-case wave-induced loads affecting the suspended payload. Motion reduction is explored through both passive hoist-arrangement optimization (geometric detuning) and active tugger control, where horizontal forces are applied to stabilize the jacket. Idealized actuator and sensor dynamics are used for the linear synthesis; since the current system lacks physical tuggers, adopting a simplified model is appropriate for this stage. A detailed structural assessment of hoist wire tensions and internal jacket stresses induced by the tugger forces falls outside the scope of this work; it is fundamentally assumed that mitigating the relative sway motion inherently minimizes dynamic tension variations across the suspension system.

1.5 Research Questions

To reach the research objective, the following main research question is defined:

> *To what extent can a control-oriented reduced-order model of the suspended jacket, combined with hoist-arrangement optimisation and tugger control, increase the operational workability envelope during JLS transit?*

This main question is answered through three sub-questions:

1. How can a reduced-order model be constructed that captures the dominant swell-induced dynamics of the suspended jacket with sufficient fidelity for control design?
2. How do hoist-arrangement parameters influence the dominant natural periods of the system, and to what extent can they be tuned passively to reduce the sway response?
3. To what extent can an active tugger-line feedback controller increase workability beyond what the hoist adjustment already achieves?

Sub-question 1 is addressed in Chapters 3 and 4, sub-question 2 in Chapter 5, and sub-question 3 in Chapter 6.

1.6 Research Theory and Methods

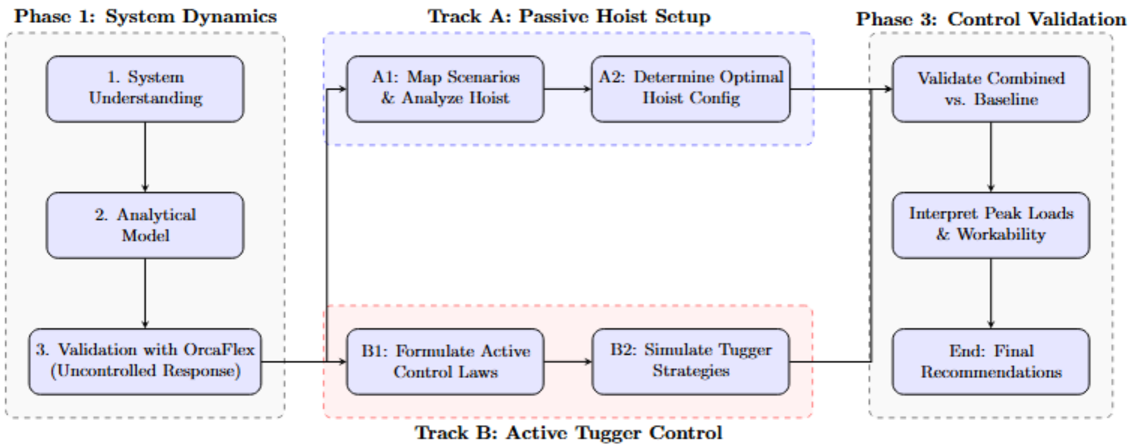

Figure 1.2: Research Method

The methodology proceeds through three sequential modelling phases, followed by two parallel branches for system improvement.

The research begins with a system-understanding phase, where the OrcaFlex baseline model is studied to establish the dominant dynamics. Response-amplitude operators (RAOs) at the hinge centroid are analyzed to identify operationally critical resonances, alongside a modal analysis of the linearized configuration to extract natural periods and mode shapes. This evaluation helps to identify the jacket sway resonance periods in model.

To facilitate rapid control synthesis, a four-degree-of-freedom (4-DOF) reduced-order model (ROM) is constructed analytically. The generalized coordinates $\mathbf{q} = [y, z, \phi, \psi]^\mathsf{T}$ capture the critical sway, heave, roll, and yaw motions. The equations of motion are derived using Lagrangian mechanics ($\mathcal{L} = T - V$). By formulating the kinetic and potential energies in closed form and applying the Euler–Lagrange equations linearized about static equilibrium, the system is reduced to a second-order linear time-invariant form:

$$\mathbf{M}\,\ddot{\mathbf{q}} + \mathbf{C}\,\dot{\mathbf{q}} + \mathbf{K}\,\mathbf{q} = -\mathbf{K}_b\,\mathbf{q}_b(t)$$

Here, the boundary excitation $\mathbf{q}_b(t)$ is derived directly from the vessel RAOs via rigid-body kinematic transfer.

To ensure the reliability of this linear ROM, it is validated against the fully nonlinear OrcaFlex reference through different tests. First, free-decay tests are used to simultaneously validate the analytical mass and stiffness matrices while calibrating the global damping ratio. Second, harmonic excitations different periods are applied to validate the sway and yaw transfer functions. Finally, the system is subjected to irregular JONSWAP wave excitation to verify the 4-DOF model validity.

The subsequent stages of the research split into two parallel branches targeting motion reduction. The first branch focuses on passive hoist optimisation, utilizing the validated ROM to sweep the design space (e.g., wire spread, anchor positions, pretensions) and rank configurations by their passive workability. The second branch targets active tugger control, where a closed-loop feedback law is synthesized on the linear ROM. The tuggers are later combined with the optimized hoist configuration also to see the combined benefits.

The system-understanding, model-building and model-validation phases are presented in Chapters 3 and 4. The hoist-arrangement optimisation and the tugger-control study follow in Chapters 5 and 6, and are compared against each other at the end of Chapter 6.

1.7 Report Structure

The remainder of this report is organised as follows.

Chapter 2 covers both the governing physics of the suspended jacket and a review of the existing motion-control literature, and closes by identifying the research gap.

Chapter 3 describes the Pioneering Spirit, the JLS, and the jacket under study, and develops the reduced-order model: the generalised coordinates, kinematic transformations, kinetic and potential energies, Euler–Lagrange derivation, and linearisation are all set out, and the modal analysis of the resulting linear time-invariant system is presented.

Chapter 4 validates the linear model against the OrcaFlex reference. Free-decay, harmonic and irregular-wave tests are presented in turn, followed by a comparison at identical excitation and a final check against a measured transit record. Together with Chapter 3 this answers sub-question 1 and establishes the model.

Chapter 5 uses the validated model to optimise the rigging geometry, sweeping the hoist parameters, identifying the detuned configuration and verifying it in OrcaFlex. This chapter answers sub-question 2.

Chapter 6 formulates and tunes an active tugger controller, reports what it achieves across the design envelope and at the edge of the operating window, and compares it against the optimised rigging on a common basis. This chapter answers sub-question 3.

Chapter 7 answers the research questions directly and gives recommendations for further work.

Chapter 2

Theoretical Background and Literature Review

This chapter establishes the physical theories and reviews the literature required to construct, reduce, and control the numerical model of the suspended jacket. It first traces the physical propagation of environmental loads starting from the sea state, through the vessel hydrodynamics, and across the extended lever arms into the suspended payload and hoist wires. It then reviews current methodologies for control-oriented reduced-order modelling and evaluates existing active and passive motion-control strategies. Finally, the chapter defines the operational workability criteria and synthesises the literature to identify the specific research gaps addressed in this thesis.

2.1 Wave Loading and Sea State Formulation

The dynamic excitation of an offshore lifting system in a heavy-lift vessel like Pioneering Spirit originates primarily from the waves acting on the vessel. Because individual sea states are irregular, they are described statistically by a wave energy spectrum $S(\omega)$, which distributes the variance of the sea-surface elevation over frequency [1]. The sea state is summarised by its spectral moments: the zeroth moment yields the significant wave height H_s, and the ratios of higher moments define characteristic periods such as the zero-up-crossing period T_z [1].

For fetch-limited or developing seas, and particularly for the swell-dominated environments critical to resonance-sensitive marine operations, the JONSWAP spectrum provides the standard parametric representation [2]. The JONSWAP spectrum is formed by multiplying the fully developed Pierson–Moskowitz spectrum by a peak-enhancement factor γ, which sharpens the spectral energy around the peak frequency ω_p:

$$S_{\mathrm{J}}(\omega) = S_{\mathrm{PM}}(\omega)\,\gamma^{\exp\left[-\dfrac{(\omega-\omega_p)^2}{2\sigma^2\omega_p^2}\right]}, \qquad \sigma = \begin{cases}\sigma_a, & \omega \le \omega_p,\\ \sigma_b, & \omega > \omega_p.\end{cases} \tag{2.1}$$

Lucas and Guedes Soares [3] report that JONSWAP reproduces the significant wave height and spectral shape of measured swell more accurately than alternative distributions.

The environment also includes current and wind, both of which are less significant compared to wave induced vessel motion transfered through the beams. Current acts mainly on the submerged hull and is counteracted by the dynamic-positioning system, making its effect on payload motion negligible; wind loads are small relative to the vessel excitation loads and may be neglected for a simplified linear model [1, 4].

2.2 Vessel Motion and Kinematic Transfer

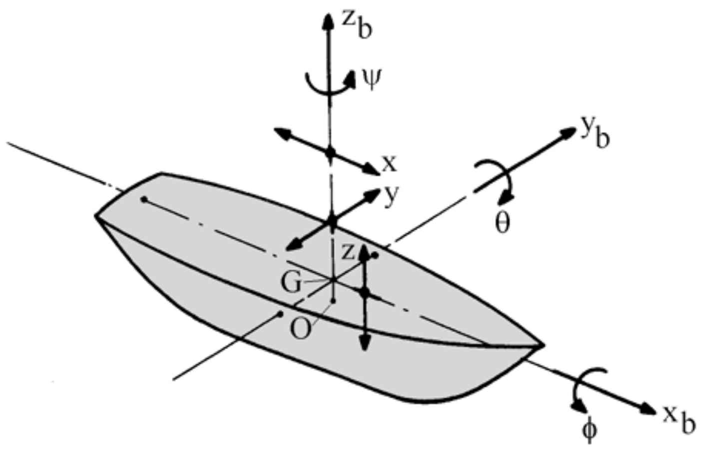

Figure 2.1: Definition of ship motions in six degrees of freedom.

The incident wave field excites the vessel as a rigid body in six degrees of freedom: surge, sway, heave, roll, pitch, and yaw. The linear transfer from wave elevation to vessel motion in each degree of freedom is quantified by Response Amplitude Operators (RAOs), which, combined with the wave spectrum, yields the statistical vessel motion response for a given sea state [1].

The excitation driving a suspended offshore payload, however, is rarely applied at the vessel centre of gravity (CoG). For stern-mounted lifting frames, the suspension hinges are located at a significant distance from the CoG. Because the lifting structure is stiff relative to the hoist wires, the crane tip is treated as a rigid offset [5]. The vessel RAOs are therefore transferred to the crane tip kinematically; processing CoG RAOs into a crane-tip excitation for a suspended-load model is the standard approach, as demonstrated by Mackojc and Chilinski [6].

For small angular motions, the vertical crane-tip displacement η_{ct} superposes vessel heave η_3 with the roll η_4 and pitch η_5 contributions through the transverse offset b and longitudinal offset l:

$$\eta_{ct} \approx \eta_3 + b\,\eta_4 + l\,\eta_5, \tag{2.2}$$

with rotations in radians. For heavy-lift vessels featuring substantial lever arms, the rotational terms dominate the kinematic transfer. A marginal vessel roll or pitch is amplified into significant translation at the suspension point, meaning the dominant boundary excitation of the payload often arises from the rotational vessel modes rather than pure heave.

2.3 Dynamics of the Suspended Payload

Because the masses of jackets that are usually transported are small relative to the displacement of the Pioneering Spirit, the reaction of the swinging jacket does not significantly alter the motion of the vessel. According to the light lift criteria established in DNV GL-RP-N103, payloads weighing less than 1-2 percent of a vessel's displacement do not alter the vessel's motion characteristics. Therefore, for such configurations, the payload motion is majorly driven by the prescribed wave-induced vessel motion. The system is therefore modelled as a base-excited rigid body hanging in air, with vessel motion as an independent boundary input.

The rigging configuration fundamentally shapes the response. Standard offshore lifts in typical heavy lift vessels which lifts monopiles or turbines often uses parallel bi-filar or single wire rigging that decouples swing (sway) from twist (yaw); the Jacket Lift System instead uses a four-wire rigging to secure the spatial footprint of the jacket. This introduces off-diagonal terms in the restoring stiffness matrix $\mathbf{K}$. Setting damping and forcing to zero, the generalised eigenvalue problem is

$$\left[\mathbf{K} - \omega^2 \mathbf{M}\right] \hat{\boldsymbol{q}} = \mathbf{0}, \tag{2.3}$$

with $\mathbf{M}$ the mass matrix and $\hat{\boldsymbol{q}}$ the mode shapes. The off-diagonal terms make the natural modes cross-coupled: a unidirectional hinge disturbance (pure sway in beam seas) forces a multi-directional response, twisting the jacket in yaw. Multi-cable suspensions are known to introduce inter-axis coupling of this kind [7]. The primary transit hazard arises when the wave period approaches the natural periods of the jacket or vessel's natural roll periods, causing resonant amplification.The vessel's RAOs act as a frequency filter for the excitations reaching the jacket.The resulting jacket motions during the transport phase is resultant of the combination of the waves hitting the vessel and the inherent dynamics of the vessel and of the jacket suspended in JLS.

2.4 Hoist Wire Tension and Structural Coupling

In the Jacket Lift System (JLS), four hoist wires suspend the jacket; the wire arrangement on the *Pioneering Spirit* is detailed in Chapter 3. The hoist wires are the mechanical links transmitting the static weight and the dynamic reaction forces between vessel and payload. For a payload suspended in air, two stiffness mechanisms act, each governing a different class of motion.

The first is the geometric restoring stiffness of pendular motion. When the payload displaces laterally, gravity provides the restoring force; for small swing angles this is equivalent to a lateral spring

$$k_s = \frac{W}{h}, \tag{2.4}$$

with W the payload weight and h the drop height from suspension point to centre of gravity. This stiffness is independent of the wire material and governs the low-frequency sway and yaw pendulum modes.

The second is the axial elastic stiffness of the wire itself,

$$k_e = \frac{EA}{L_0}, \tag{2.5}$$

with EA the axial rigidity and L_0 the unstretched length; segments of different material in a single line combine in series. This stiffness governs the vertical (heave) response of the payload relative to the suspension point.

Because $k_e \gg k_s$ in practical heavy-lift rigging, the two mechanisms are well separated: the axial natural period lies well below the dominant wave periods, while the pendulum natural periods lie well above them. There are some studies which focused on the heave compensations [4, 8] where wave-frequency excitation acts directly along the wire axis. But during suspended transit the axial mode responds quasi-statically and dynamic amplification in pure heave is small. The severe cyclic tension fluctuations are instead kinematically driven by the low-frequency pendulum motions. As the payload swings the specific hoist geometry forces the lines to extend and contract diagonally, redistributing load among the lines and inducing tension peaks even at modest sway amplitudes, as illustrated in Figure 2.2. This pendulum-driven tension mechanism which is a result of the specific four hoist rigging arrangement characterize the risks of sway motion.

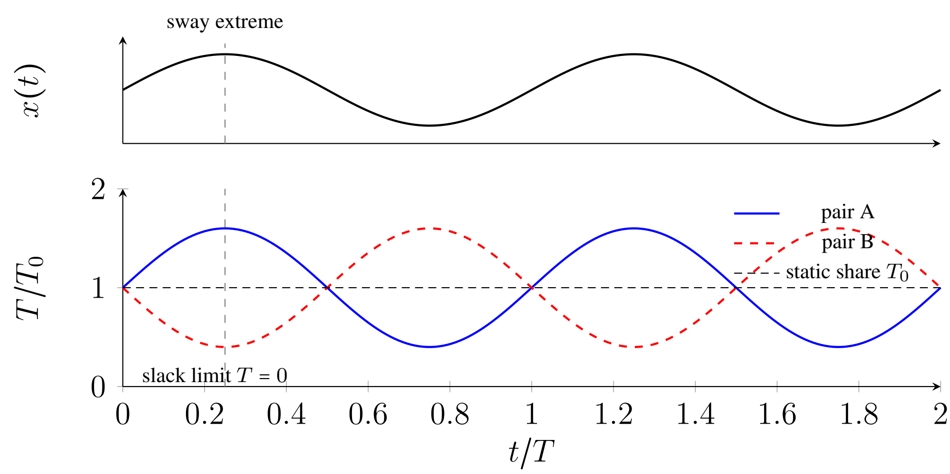

Figure 2.2: Kinematic tension redistribution during pendular sway (schematic). Each crossed diagonal wire pair is lumped as a single equivalent axial spring. The pair tensions oscillate in anti-phase about the static share T_0 (bottom).

In subsea lifts, hydrodynamic drag on the submerged rigging and payload is a dominant dissipation mechanism (DNV-RP-N103). During transit in air, the corresponding aerodynamic dissipation is reduced by approximately the density ratio $\rho_{\text{air}}/\rho_{\text{water}} \approx 1.2 \times 10^{-3}$ at equal kinematics. Linearising the quadratic drag force gives an equivalent damping. Aerodynamic damping is therefore omitted from the linear design plant; since this removes a dissipative term, the omission is conservative for resonance-driven response. Hoist lines in air are therefore commonly idealised as undamped linear axial springs. For synthetic fibre segments such as HMPE, this idealisation requires justification: the material exhibits viscoelastic and viscoplastic behaviour, but full-scale experiments on offshore lifting slings show that a pre-operational bedding-in sequence removes non-reversible elongation, after which the dynamic axial stiffness stabilises under cyclic loading [9]. A properly bedded-in HMPE segment can thus be represented as a linear spring in series with the steel wire.

2.5 Control-Oriented Reduced-Order Modelling

High-fidelity nonlinear time-domain tools (e.g. OrcaFlex) resolve the spatial wire geometry and instantaneous tensions accurately, but a single representative sea-state run is computationally expensive [10, 11]. The research questions of this thesis — sweeping hoist-arrangement parameters and finding a tugger control strategy — require many model evaluations, which is infeasible with the full nonlinear reference [10, 11]. This motivates a *control-oriented reduced-order model*: a low-order, linear time-invariant description retaining only the dominant resonant dynamics, cheap enough for high-speed iteration.

The literature presents two distinct analytical approaches for offshore lifting: fully coupled payload–vessel models and uncoupled base-excited models. Coupled models capture the mutual dynamic interactions between the swinging payload and the vessel's hydrodynamics [6]. Conversely, uncoupled models treat the vessel motion as a prescribed boundary input, utilizing Response Amplitude Operators (RAOs) to derive the crane tip excitation [6, 12]. The limit of applicability between these approaches depends strictly on the payload-to-vessel mass ratio [6]. Mackojc and Chilinski [6] explicitly demonstrate that uncoupled, RAO-driven models remain dynamically valid for "light lifts" where the payload is less than 1–2% of the vessel's displacement; beyond this threshold, the payload significantly alters the vessel's inherent motion characteristics, necessitating a coupled analysis.

For operations fitting the light-lift criteria, uncoupled analytical models offer a low computational cost while still successfully capturing relevant physical phenomena [12]. Chilinski et al. [12] highlight that while many simplified models uncouple heave and pendulation entirely, a properly formulated base-excited model can accurately detect parametric resonance and evaluate compensating elements. Recent control-oriented studies heavily rely on this uncoupled methodology to overcome the prohibitive computational times of full 3D simulations. Eerden [10] employs a 2D analytical approach, modelling the crane and hoist as a damped elastic pendulum driven by predefined vessel trajectories, to design and tune tugger control systems. Similarly, Sanders [11] utilizes a reduced-degree-of-freedom mass-spring-damper model, with equations of motion derived via Lagrangian mechanics, to rapidly iterate through controller tunings.

In both cases, the simplify-then-validate-against-a-high-fidelity-reference methodology is standard practice for offshore lifting [10, 11]. Sanders [11] and Eerden [10] successfully validate their reduced analytical models against OrcaFlex before proceeding with control design. Because the *Pioneering Spirit* operates with a massive displacement relative to the transported jackets, the base-excited assumption holds [13]. The resulting linear plant lets vessel motion enter as a prescribed boundary input, providing the computationally efficient foundation required to rapidly evaluate both passive geometric detuning and active, velocity-proportional control laws.

2.6 Motion-Control Strategies for Suspended Payloads

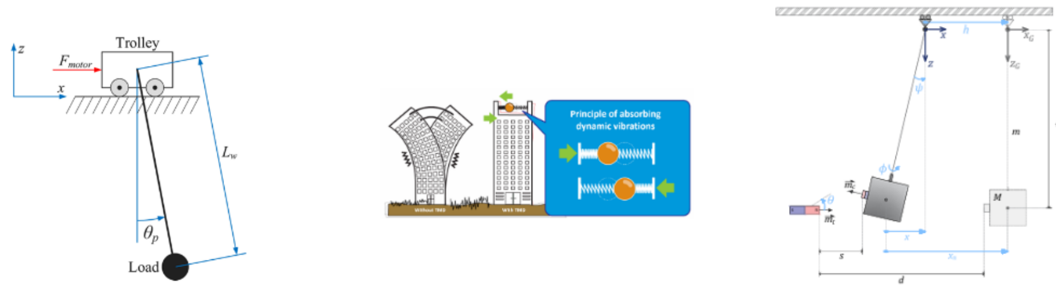

Figure 2.3: Overview of motion-control concepts for suspended offshore payloads:anti-swing slider (left, adapted from Jensen et al. [14]), tuned mass damper principle (centre, adapted from Makita [15]), and contactless electromagnetic control (right, adapted from Atzampou [16]).

Figure 2.3 illustrates three representative concepts currently being studied in the industry. Anti-swing slider systems (left) suppress pendulation by translating the suspension point itself; at heavy-lift scale the required actuator forces and the effect of the moving mass on crane capacity are prohibitive, and the JLS suspension points are in any case fixed at the TLB tips [4]. Tuned mass dampers (centre) absorb energy through an auxiliary oscillator but several practical reasons makes it less attractive in offshore operations . Contactless electromagnetic actuation (right) can exert both attractive and repulsive forces with a force-to-weight ratio comparable to active tugger lines, but remains at the proof-of-concept stage, validated on a laboratory-scale pendulum and requiring magnetic attachments on the payload [17].

Controlling a suspended offshore payload is complicated by its underactuated nature, having fewer actuators than degrees of freedom, and the continuous motion of the crane base [18]. The strategies relevant here fall into two categories: *active* control, which dissipates energy through commanded actuator forces, and *passive* control, which alters the physical dynamics so that resonance is avoided. They both can be complementary as the passive detuning lowers the peak gain an active controller must handle, and active control covers the residual response detuning cannot remove. The broader landscape of motion-reduction hardware concepts—such as anti-swing slider systems, heave compensation, gyroscopic stabilisation, and automated side loaders—has been evaluated and ranked elsewhere through multi-criteria feasibility studies [4]. Gyroscopic stabilisation relies on conservation of angular momentum, and the required gyroscope mass and spin rate scale with the payload, making it impractical at this scale. Because this thesis is constrained to the operational architecture of the JLS, it focuses on tugger-line control and geometric detuning: the active and passive avenues natively applicable to this setup.

2.6.1 Tugger Lines

To mitigate wave-induced pendular sway of the suspended jacket during transit without modifying the vessel hull or primary hoist structure, active tugger winches apply dynamic control forces at the payload base [10, 19].

Evolution of Control Modes: Fixed, Constant Tension, and Active Damping
Three distinct operational modes for tugger lines are documented in heavy-lift literature:

1. **Fixed Tugger Lines:** Modelling tugger wires as fixed-length cables introduces an effective horizontal spring stiffness to the system [20, 21]. Adding this fixed horizontal stiffness lowers the natural pendular period of the payload, potentially shifting it directly into the peak wave energy band of the sea state [11]. This forced shift can lead to severe resonance excitation, decreased workability, and an increased risk of wire overloads or catenary snap loads [11, 19].

2. **Constant Tension (CT) Mode:** In this mode, the winch attempts to maintain a fixed tension setpoint regardless of payload movement. Because the theoretical tension-velocity slope is zero ($\partial F/\partial v = 0$), CT mode provides no active velocity damping, allowing un-damped pendular sway near resonance [10]. While physical imperfections in winch actuation may inadvertently dissipate a small amount of energy in practice, the mode itself does not actively mitigate pendulation [10, 11].

3. **Active Damping Mode:** Introduced by Meskers and Dijk [19], active damping tuggers dynamically adjust line tension as a function of payout velocity. By commanding a higher tension during pay-out ($v > 0$) and a lower tension during pay-in ($v < 0$), energy is continuously extracted from the swinging payload over each cycle. This effectively provides active damping without undesirably altering the system's natural stiffness [11, 19].

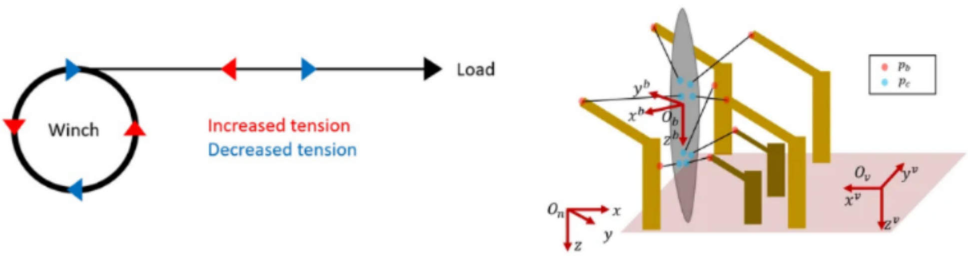

Figure 2.4: Schematic diagram of an active tugger line (left, adapted from Eerden [10]) and a multi-cable arrangement (right, adapted from Ren et al. [22]).

Setpoint Laws and Sensor Topologies
The effectiveness of an active damping tugger is fundamentally governed by its setpoint calculation law $f(v)$, which translates sensed motion into a tension demand. Early industrial implementations relied on piecewise algebraic profiles, mapping payload velocity to tension through one-dimensional lookup curves bounded by the wire's minimum and maximum safe working loads [19, 20]. Evaluated profiles in the literature include linear, stepwise (bang-bang), and quadratic laws. However, these algebraic profiles possess inherent dynamic shortcomings: Sanders [11] demonstrated that quadratic profiles apply insufficient damping force at low turnaround velocities, while stepwise switching introduces force discontinuities that cause severe winch drive chatter.

To resolve these discontinuities, continuous Proportional-Derivative (PD) setpoint laws are frequently utilized. Pure proportional velocity feedback suffers from turnaround phase lag, applying maximum force too late in the swing cycle and occasionally feeding energy back into vessel–payload coupling modes [10, 11]. To counteract this, a derivative term acting on the acceleration error is introduced to provide anticipatory phase-lead [10]. The resulting control law commands a tension F_s bounded by the physical limits of the system:

$$F_s = \max\left(F_{\text{min}},\ \min\left(F_{\text{max}},\ F_{\text{pre}} + \bar{k}_p e_v + \bar{k}_d \dot{e}_v\right)\right) \tag{2.6}$$

where e_v is the velocity error, and $\bar{k}_p$ and $\bar{k}_d$ are the proportional and derivative gains. Integral action is explicitly omitted from this architecture because the continuously varying cyclic nature of wave excitation inevitably causes integrator anti-windup saturation and subsequent control instability [10].

The accuracy of this control law relies entirely on the physical location of the feedback sensor. Controllers typically measure motion either indirectly via a winch fairlead encoder or directly via on-deck hinge and payload sensing. Relying solely on winch payout assumes the wire is perfectly stiff; in reality, elastic compliance (stretching) and catenary (sagging) effects introduce uncertainties between the winch response and the actual payload kinematics, making direct payload sensors much better for accurate tracking [10].

Actuator Constraints and Capacity Scaling
Realizing any theoretical control law in a time-domain simulation requires enforcing strict actuator constraints. Rather than assuming ideal force application, the literature dictates that active damping must account for installed capacity ratios, which scale the available winch force against the payload mass [11, 19]. While exact ratios vary widely depending on the vessel and operation, all active systems are constrained by a maximum Safe Working Load to prevent wire yield, and a minimum tension floor to prevent catenary snap loads [10, 11].

Furthermore, kinematic limits on the rated line payout speed and thermal motor power caps heavily restrict energy dissipation during severe sea states. The physical inertia of the winch drum and the first-order response lag of the motor drive introduce mechanical delays that must be modeled to accurately predict whether the active system will successfully dampen or inadvertently amplify the payload resonance [10].

2.6.2 Passive Control: Geometric Detuning and Tuned Mass Dampers

Passive strategies alter the physical dynamics of a system to avoid excitation without requiring external energy or active actuators. In the offshore industry, Tuned Mass Dampers (TMDs) are a standard passive solution for suppressing crane and structural vibrations [23, 24]. A TMD absorbs energy through an auxiliary oscillator tuned to the problematic resonance frequency. However, for heavy-lift transits involving payloads of several thousand tonnes, the required auxiliary mass becomes prohibitively large, imposing an unacceptable penalty on the vessel's overall lift capacity.

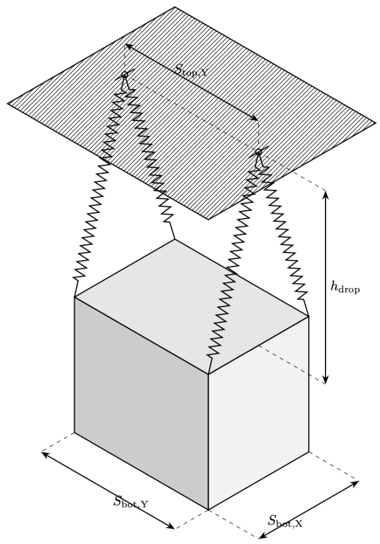

Figure 2.5: 3D Isometric View

A mass-efficient alternative for suspended payloads is *geometric detuning*. In a simple mathematical pendulum, the natural period is governed exclusively by its length; increasing the drop height physically lowers the natural frequency. The Jacket Lift System, however, is a complex spatial arrangement suspended by four separate hoist wires. By systematically altering these specific hoist-arrangement parameters—such as the top wire spread, bottom anchor positions, and overall drop height—the geometric restoring stiffness (Equation 2.4) of the entire multi-wire configuration is modified. This allows the system's dominant natural periods to be deliberately shifted out of the hazardous bands, avoiding the resonance trap without active mechanical intervention.

The necessity for this detuning is sharpest when operational parameters force a lift to occur near its natural period. As highlighted by McKenzie and Irani [25], operating at a pendulum's natural frequency represents a worst-case scenario, creating a dangerous condition where resonant amplification is bounded solely by inherent structural damping or minimal aerodynamic drag. While the standard approach to mitigating this severe resonance relies on complex active feedback and compensation frameworks [25], systematic geometric detuning offers a purely passive, preemptive solution. For a spatial, multi-point heavy-lift payload of the scale examined in this thesis, systematic detuning via rigging geometry remains unaddressed in the existing literature; this gap is the focus of the subsequent passive hoist optimisation.

2.7 Operational Limits and Workability Criteria

The engineering objective of motion control is to widen the operational weather window, quantified through workability limits. Guachamin Acero et al. [26] established that assessing marine operability requires identifying critical failure events — for JLS transit, peak liftwire tensions exceeding the Safe Working Load and exhaustion of clearance between the swinging jacket and the TLB structure.

For a linearised system, the structural response is proportional to H_s at fixed T_z. For offshore operations, DNV-RP-H103 [5] computes the particular response of the system X at unit wave height ($H_s = 1\,\mathrm{m}$) and scales it against the design limit $X_{\max}$ to set the allowable sea-state boundary:

$$H_{s,\mathrm{lim}}(T_z) = \frac{X_{\max}}{X(H_s = 1\,\mathrm{m},\, T_z)}. \tag{2.7}$$

In practice, offshore heavy-lift operations are governed by multiple simultaneous criteria, and the overall workability limit is dictated by whichever constraint yields the lowest allowable wave height $H_{s,\mathrm{lim}}(T_z)$. Crucially, the introduction of a motion control system can fundamentally shift this governing bottleneck.

As demonstrated by Meskers and Dijk [19], successfully applying a damping tugger system to suppress transverse payload pendulation can improve the operability of that specific state so significantly that it is no longer the limiting condition for the entire operation. Consequently, when evaluating control and detuning strategies, it is essential to concurrently monitor all dynamic responses; mitigating the primary failure mode may successfully extend the weather window, but it will inevitably transfer the operational limit to a different criterion (e.g., the transverse motion criterion, longitudinal motion criterion, vertical motion criterion, side lead criterion, or hoist wire DAF criterion). Operability is then the statistical probability that the site sea state stays below this newly established boundary.

2.8 Research Gap and Conclusion

While the individual physical components of this system—namely wave excitation (Section 2.1), kinematic transfer across the extended lever arm (Section 2.2), base-excited payload dynamics (Section 2.3), and motion control (Section 2.6)—are well-established, three distinct gaps persist regarding the *Pioneering Spirit* Jacket Lift System (JLS).

First, current payload-dynamics and control literature typically assumes parallel one-hoist or two-hoist rigging arrangements that result in standard vertical pendulation. To the author's knowledge, these frameworks fail to capture the specific four-hoist wire arrangement that drives the critical sway-to-yaw coupling inherent to this operation.

Second, while generic reduced-order offshore-lift models are available [6, 12], and a control-oriented linear model of the specific JLS hoist system exists for the lowering phase of unloaded hoists [27], neither captures the cross-coupled stiffness of the loaded four-wire transit configuration. This renders them unsuited for the rapid hoist-parameter and controller iterations demanded during early design phases.

Third, passive resonance mitigation in offshore lifting usually relies heavily on Tuned Mass Dampers (TMDs) [23, 24]. The alternative concept of geometric detuning via optimised hoist configuration remains an unexplored area for spatial payloads of this magnitude.

These gaps present a practical modelling problem. High-fidelity nonlinear software accurately captures the complex JLS geometry but is too computationally expensive for iterative design sweeps and controller tuning. To resolve this issue, this thesis details the formulation and validation of a computationally efficient, control-oriented reduced-order model tailored specifically to this four-hoist configuration. This model is subsequently deployed to evaluate both passive geometric detuning and the feasibility of active tugger line control strategies, bridging the operational and analytical gaps identified above.

Chapter 3

System Description and Model

This chapter establishes the physical parameters of the Jacket Lift System (JLS) and formulates the control-oriented model used for iterative design. It first defines the physical lifting setup and the high-fidelity ORCAFLEX reference model. Subsequently, a 4-DOF analytical model is derived in MATLAB. The underlying physical simplifications, modal verification, and damping identification required to construct this linearised model are detailed, concluding with a summary of the finalised system parameters.

3.1 The Physical System and Lifting Setup

Figure 3.1: Pioneering Spirit

The *Pioneering Spirit* is a twin-hulled heavy-lift vessel equipped at its stern with the Jacket Lift System (JLS). The JLS carries the jacket on two Tilting Lift Beams (TLBs)—long arms hinged at the stern to support the hoisting arrangement at their tip. During the transit phase, the jacket hangs vertically from the TLB tips while the vessel transits a seaway; this suspended-transport condition governs the dynamic loads of interest.

Figure 3.2: An offshore jacket handled during heavy-lift installation.

For the modelling, the base-case payload of a steel jacket of dry mass 7650 t is chosen . Jackets can weigh 2000 to 12000 tonnes in current projects. This payload is suspended in air from four main hoist (MH) wires. The wires are arranged in an inverted V-shape longitudinally and a slightly inclined vertical parallelogram configuration transversely: they run from two pairs of top anchors on the beams—an inner pair connected to padeyes near the vessel, and an outer pair connected to padeyes further outboard—totalling four nearly symmetric lift points on the jacket.

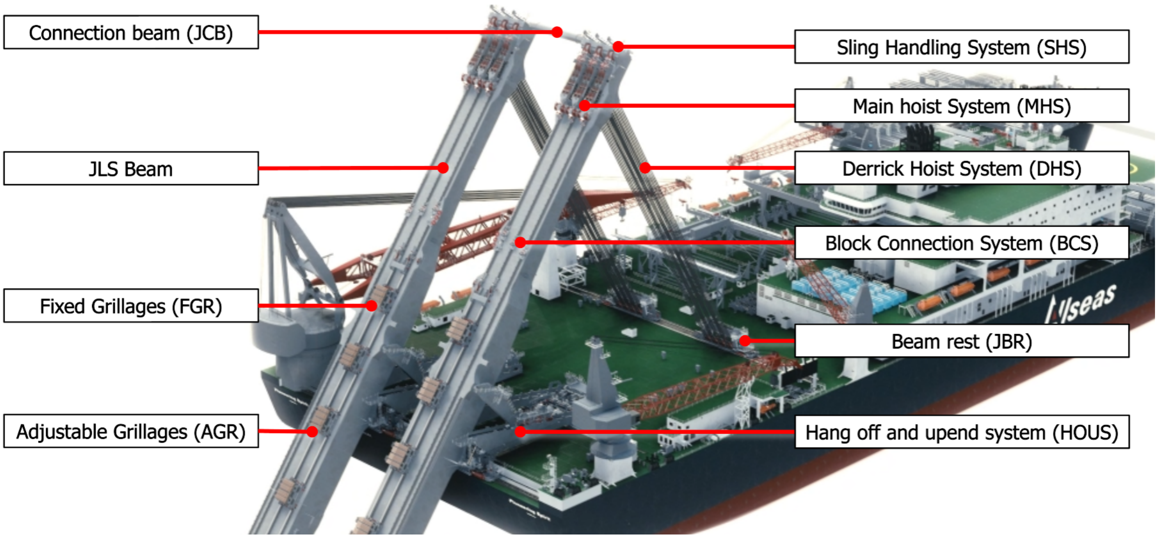

Figure 3.3: Overview of the primary structural and mechanical components comprising the Jacket Lift System (JLS) on the Pioneering Spirit.

Figure 3.4: BCS

Additionally, the physical hoist setup includes 190 t mass blocks attached to the wire assemblies to maintain tension, and a Block Connection System (BCS) comprising constant-tension winches and nonlinear damper links connecting the hoist blocks to the lift beams to restrict motion during lowering and grommet connection phases.

3.2 The OrcaFlex Reference Model

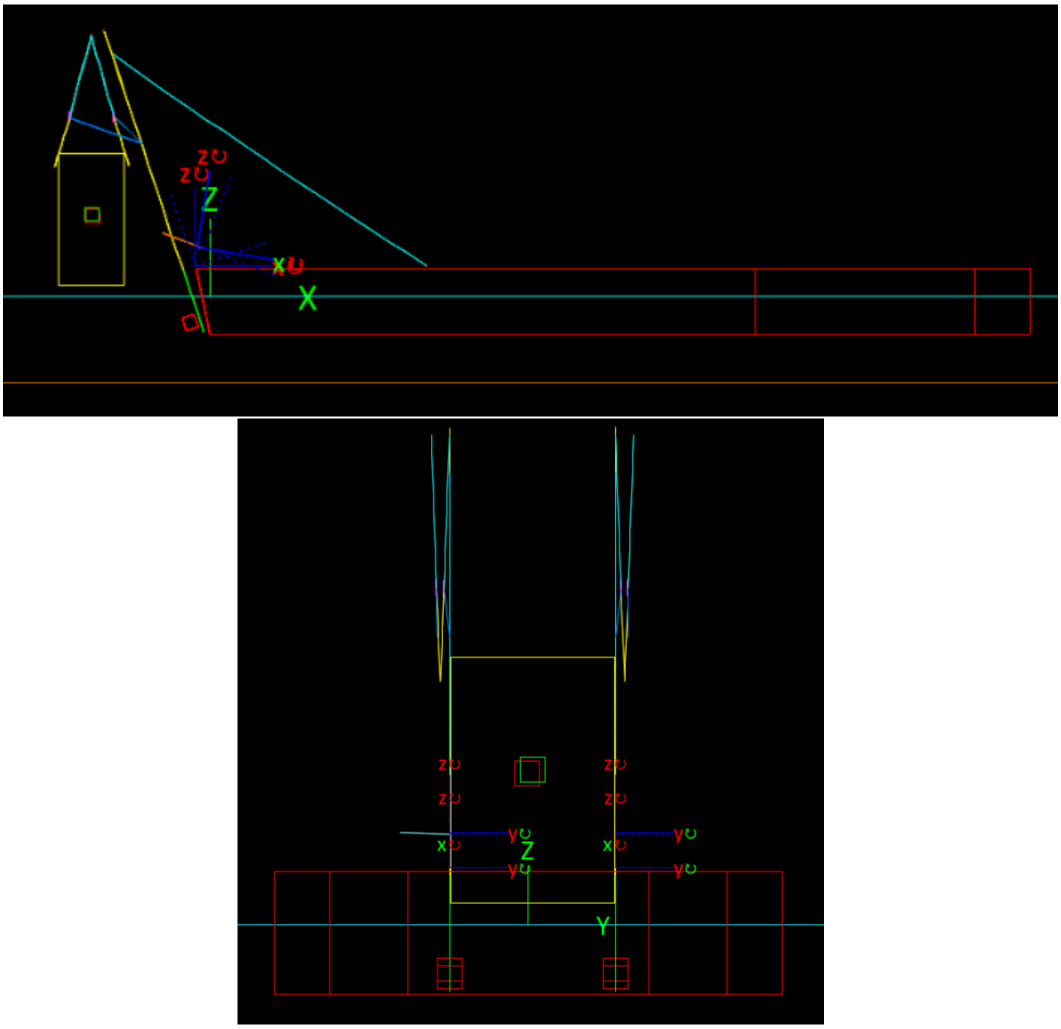

Figure 3.5: Visualisation of the physical system and ORCAFLEX reference model showing the *Pioneering Spirit*, the JLS beams, and the representative suspended jacket.

A high-fidelity, fully non-linear time-domain model of the physical system constructed in ORCAFLEX was used to serve as validation reference. In this model, the jacket is represented as a rigid body with accurately distributed mass and inertia. The four hoist wires are modelled as finite-element axial springs capable of capturing exact spatial geometry and dynamic tensions.The model explicity uses wires of two properties namely MH and Dyneema to model the hoist wires.The MH blocks which sits between hoist wire ends is also modeled.The Orcaflex models includes the whole vessel modeled with RAOs and detailed connections of the Jacket and JLS. While this model captures the full nonlinear behaviour of the transit operation, its high computational cost makes it impractical for evaluating extensive sea states. This requires the development of a control-oriented reduced model. The Orcaflex model includes the BCS, main hoist wires, and the MH mass blocks, which are modelled with greater complexity.The BCS is modelled through winches and links, incorporating specific material definitions for the distinct wire segments.The Jacket payload is modeled using 6D bouys with drag areas,Mass and Inertia.

3.2.1 Kinematic Excitation from Waves

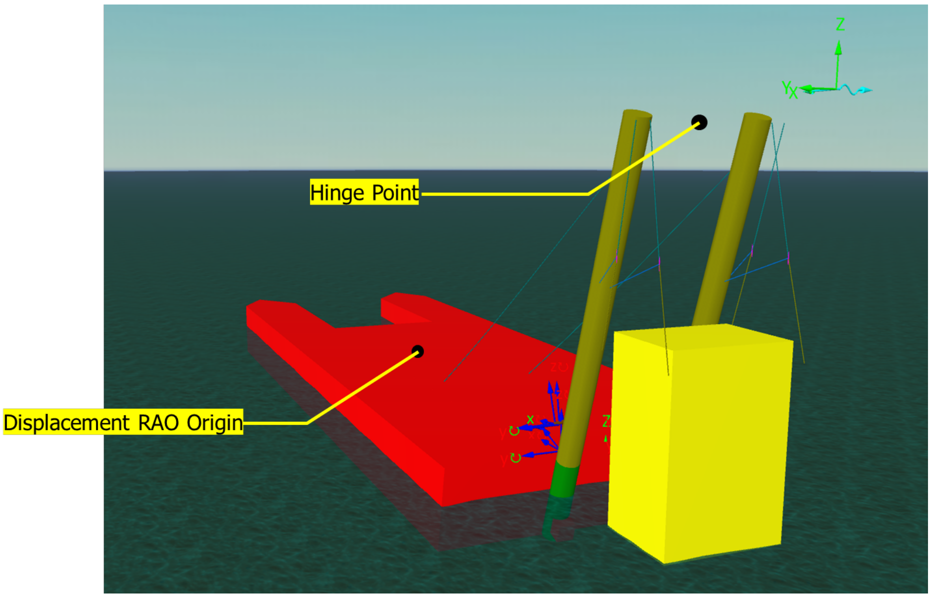

Figure 3.6: Definition of the kinematic lever arm separating the vessel's displacement RAO origin and the TLB hinge point within the OrcaFlex environment.

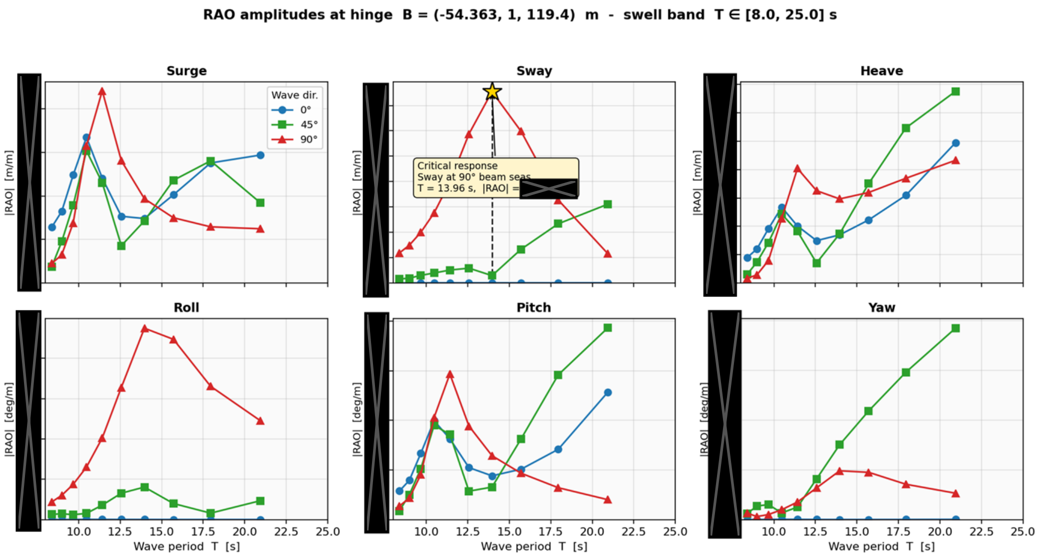

Figure 3.7: Displacement RAOs transferred from the vessel CoG to the TLB hinge.

Within the OrcaFlex environment, vessel displacement RAOs are inherently referenced to the vessel's RAO origin. However, the boundary excitation driving the jacket is applied at the TLB hinge in ROM. To map the ORCAFLEX reference to the analytical MATLAB model, the vessel hydrodynamics are simplified into a prescribed boundary input defined by the RAO-based hinge excitations. This motion must be transferred via rigid-body kinematics over the substantial lever-arm vector $\mathbf{r}_{B/CoG} = (-215.263,\ 1.0,\ 99.998)\ \mathrm{m}$ between the vessel CoG (O) and the hinge centroid (B). For the small vessel angles typical of transit, the linearised transfer relations defining the hinge input ($\mathbf{q}_b$) are:

$$\begin{aligned} H_{\text{sway,B}}(\omega) &= H_{\text{sway,O}}(\omega) + H_{\text{yaw,O}}(\omega)\,\Delta x - H_{\text{roll,O}}(\omega)\,\Delta z, \\ H_{\text{heave,B}}(\omega) &= H_{\text{heave,O}}(\omega) + H_{\text{roll,O}}(\omega)\,\Delta y - H_{\text{pitch,O}}(\omega)\,\Delta x, \\ H_{\text{roll,B}}(\omega) &= H_{\text{roll,O}}(\omega), \\ H_{\text{yaw,B}}(\omega) &= H_{\text{yaw,O}}(\omega), \end{aligned} \tag{3.1}$$

where $(\Delta x, \Delta y, \Delta z) = \mathbf{r}_{\text{B/CoG}}$. Due to the extreme longitudinal offset ($\Delta x \approx -215\,\text{m}$), the vessel's rotational motions (specifically roll and pitch) are amplified into large translational motions at the hinge. This kinematic amplification dictates the dominant boundary conditions driving the payload, enabling the MATLAB model to function entirely on these isolated hinge displacement RAOs.

3.3 Development of the MATLAB Model

To facilitate the high-volume parametric sweeps required for hoist optimisation and iterative controller tuning, a simplified, frequency-domain analytical 4-DOF model was formulated in MATLAB. The system kinematics are defined using a fixed inertial Earth frame alongside two moving body-fixed frames: one located at the vessel's centre of gravity (CoG) and one at the jacket's CoG.

3.3.1 Modelling Assumptions

Deriving the linear state-space model (Equation 3.10) requires several physical simplifications and parameter reductions.

Equivalent Series Stiffness of the Hoist Rigging

The main hoist lines are not uniform; they consist of a steel wire rope segment connected to a synthetic Dyneema rope segment, with a heavy MH block positioned at their physical intersection.

While the mass of the MH block is dynamically reallocated to the jacket's CoG via kinetic energy equivalence, the two cable segments act strictly in series structurally. As a result, they experience the same axial force, and their total dynamic deflection is the sum of their individual deflections. The equivalent axial stiffness (k_{eq}) of each two-part hoist line is calculated using the reciprocal sum of the individual segment stiffnesses:

$$k_{\text{eq}} = \left(\frac{1}{k_{\text{wire}}} + \frac{1}{k_{\text{dyneema}}} \right)^{-1} = \left(\frac{L_{\text{wire}}}{EA_{\text{wire}}} + \frac{L_{\text{dyneema}}}{EA_{\text{dyneema}}} \right)^{-1} \tag{3.2}$$

Given the nominal properties of the 34 m steel wire ($EA = 9.64 \times 10^6$ kN) and the 24 m Dyneema segment ($EA = 2.0 \times 10^6$ kN), the theoretical series stiffness evaluates to approximately 63 274 kN/m. The analytical model utilises this value.

Kinematic Mass Lumping (Kinetic Energy Equivalence)

The physical JLS incorporates the mass of the MH blocks, hoist wires, grommets, etc., which are not explicitly modelled in the analytical model. The steel wires are also heavy, although their mass is small compared to the jacket's mass. There are four 190 t MH blocks (totalling

760 t) located at the intersection of the steel and Dyneema segments. Because these blocks are located partway up the suspension length, they experience a smaller translational velocity than the jacket during a pendular sway motion. Simply summing the masses would artificially inflate the kinetic energy of the linear model and improperly lower the natural frequency.

Instead, to preserve the dynamic characteristics of the system within a rigid-body 4-DOF framework, the mass blocks are reduced to an effective inertial mass at the jacket CoG using kinetic energy equivalence. The mass contribution is scaled by the square of the kinematic lever-arm ratio:

$$M_{\text{kin}} = M_{\text{jacket}} + \sum m_{\text{block}} \left(\frac{s_{\text{block}}}{L_{\text{total}}} \right)^2 + \int_0^{L_{\text{total}}} \mu(s) \left(\frac{s}{L_{\text{total}}} \right)^2 ds \tag{3.3}$$

Evaluated for the baseline arrangement, this yields an effective additional mass of approximately 290 t. That estimate is an upper bound: it assumes the blocks follow the pendular velocity profile freely, whereas in the physical system the Block Connection System ties them to the lift beams and restrains that motion. Two further difficulties arise in practice. The lever-arm ratio $s_{\text{block}}/L_{\text{total}}$ is set by the hoist geometry, so the additional mass is not a fixed property of the payload but changes with every candidate arrangement examined in the optimisation study of Chapter 5; and the degree to which the BCS restrains the blocks is not known a priori. Rather than propagate an arrangement-dependent quantity of uncertain magnitude, the inertial mass is instead identified once, from the free-decay period of the bare arrangement, giving $M_{\text{kin}} = 7719$ t—an additional mass of 69 t, bracketed by the payload mass below and the free kinetic-energy estimate above. The gravitational mass ($M_{\text{grav}} = 7650$ t) is decoupled and maintained at the dry jacket weight so that the geometric pendular restoring stiffness remains accurate.

The sensitivity of the sway frequency to this choice is low, and M_{kin} is treated as a constant through the hoist-optimisation phase. The bracketing argument, the identification evidence, and the sensitivity of the sway period across the full plausible mass range are given in Appendix A.

3.3.2 Lagrangian Formulation

Based on the dominance of beam-sea excitation, the jacket is modelled as a rigid body with four degrees of freedom. Surge and pitch are excluded as jacket DOFs due to dynamic positioning constraints and the minimal longitudinal excitation in beam seas:

$$\mathbf{q} = \begin{bmatrix} y & z & \phi & \psi \end{bmatrix}^{\mathsf{T}}, \tag{3.4}$$

representing sway, heave, roll, and yaw of the Jacket at the CoG. The orientation of the jacket relative to the inertial frame is parameterised by the rotation sequence $\mathbf{R}_{\text{jacket}}(\phi, \psi) = \mathbf{R}_z(\psi)\, \mathbf{R}_x(\phi)$. The vessel orientation is parameterised identically by $\mathbf{R}_{\text{vessel}}(\phi_b, \psi_b)$. Denoting the CoG vertical offset at equilibrium as $z_{\text{eq}} = h_{\text{drop}} + z_p$, the instantaneous world-frame positions of the jacket CoG and the moving hinge centroid are:

$$\mathbf{r}_{\text{CoG}} = \begin{bmatrix} 0 \\ y \\ z + z_{\text{eq}} \end{bmatrix}, \qquad \mathbf{r}_{\text{hinge}} = \begin{bmatrix} 0 \\ y_b \\ z_b \end{bmatrix}. \tag{3.5}$$

The instantaneous length of wire i, spanning from top anchor $\mathbf{A}_i$ to bottom padeye $\mathbf{P}_i$, is defined geometrically as:

$$L_i(\mathbf{q}, \mathbf{q}_b) = \left\| (\mathbf{r}_{\text{hinge}} + \mathbf{R}_{\text{vessel}}\, \mathbf{A}_i) - (\mathbf{r}_{\text{CoG}} + \mathbf{R}_{\text{jacket}}\, \mathbf{P}_i) \right\|. \tag{3.6}$$

The kinetic energy T incorporates the effective inertial mass (M_{kin}) and is formulated as:

$$T = \tfrac{1}{2} M_{\text{kin}} \left(\dot{y}^2 + \dot{z}^2 \right) + \tfrac{1}{2} I_{xx} \dot{\phi}^2 + \tfrac{1}{2} I_{zz} \dot{\psi}^2, \tag{3.7}$$

yielding the diagonal mass matrix $\mathbf{M} = \mathrm{diag}(M_{\text{kin}},\ M_{\text{kin}},\ I_{xx},\ I_{zz})$. The potential energy V comprises gravitational energy (using static gravitational mass M_{grav}) and the strain energy of the wires, modelled as linear axial springs of stiffness k:

$$V(\mathbf{q}, \mathbf{q}_b) = -M_{\text{grav}}\, g\, (z + z_{\text{eq}}) + \tfrac{1}{2} \sum_{i=1}^{4} k \left[L_i(\mathbf{q}, \mathbf{q}_b) - L_{0,i} \right]^2. \tag{3.8}$$

Substituting T and V into the Lagrangian $\mathcal{L} = T - V$, applying the Euler–Lagrange equations, and linearising about the static equilibrium ($\mathbf{q} = \mathbf{q}_b = \mathbf{0}$) extracts the constant restoring stiffness ($\mathbf{K}$) and base-excitation ($\mathbf{K}_b$) matrices via the Hessians of the potential energy:

$$\mathbf{K} = \left. \frac{\partial^2 V}{\partial \mathbf{q}\, \partial \mathbf{q}^{\mathsf{T}}} \right|_{\mathbf{q}=\mathbf{0},\, \mathbf{q}_b=\mathbf{0}}, \qquad \mathbf{K}_b = \left. \frac{\partial^2 V}{\partial \mathbf{q}\, \partial \mathbf{q}_b^{\mathsf{T}}} \right|_{\mathbf{q}=\mathbf{0},\, \mathbf{q}_b=\mathbf{0}}. \tag{3.9}$$

Energy dissipation is represented by a viscous damping matrix $\mathbf{C}$, whose construction requires the modal solution and is deferred to Section 3.5. With $\mathbf{C}$ in place, the equations of motion are:

$$\boxed{\mathbf{M}\, \ddot{\mathbf{q}} + \mathbf{C}\, \dot{\mathbf{q}} + \mathbf{K}\, \mathbf{q} = -\mathbf{K}_b\, \mathbf{q}_b(t)\,.} \tag{3.10}$$

3.4 Modal Analysis

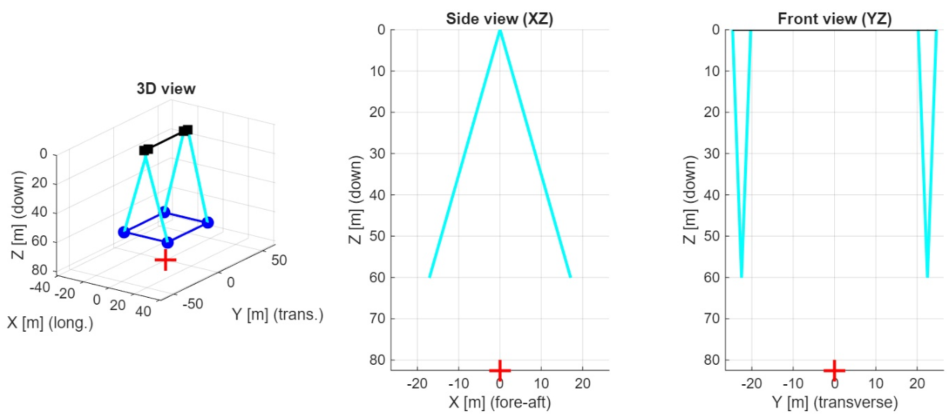

Figure 3.8: MATLAB Model Geometry Plotted

Solving the generalised eigenvalue problem confirms the analytical model accurately captures the system physics. The resulting natural periods and modal kinetic-energy fractions are reported in Table 3.1.

Mode	T_i [s]	Sway %	Heave %	Roll %	Yaw %	Dominant
1	15.7	98.7	0.0	0.0	1.3	Sway
2	3.543	1.3	0.0	0.3	98.5	Yaw
3	1.2	0.0	0.0	99.7	0.3	Roll
4	1.1	0.0	100.0	0.0	0.0	Heave

Table 3.1: Modal analysis of the linear ROM.

To validate these results, the global modes are compared against a high-fidelity 3D OrcaFlex finite element model, as presented in Table 3.2. It is important to note that the OrcaFlex model possesses a vast number of degrees of freedom, resulting in numerous higher-order eigenmodes. However, the majority of these represent high-frequency local dynamics, such as transverse wire vibrations or internal elastic deformations of the jacket structure. For the purpose of this comparative validation, the 3D model's modal results were filtered to isolate the fundamental global rigid-body modes that govern the payload's macroscopic movement.

Dominant DOF	**OFx Mode**	**OFx T_n [s]**	**ROM Mode**	**ROM T_n [s]**
Surge	1	19.7	-	-
Sway and Yaw	2	15.5	1	15.7
Yaw and Sway	3	4.4	2	3.543

Table 3.2: Modal comparison between the full 3D Orcaflex model and the linear 4-DOF MATLAB model.

The primary sway mode of the MATLAB model at $T_1 = 15.7\,\mathrm{s}$ agrees with the 3D model's second mode (15.5 s, a deviation of 1.3%) and lies at the principal resonance trap of the swell band. The yaw mode at $T_2 = 3.543$ s lies well below the swell band, as does the 3D model's third mode at 4.4 s. The reduced-order period is 19.5% shorter, but the yaw coordinate is not resonantly excited.

3.5 Damping of the Suspended Jacket in Air

The modal solution above is that of the undamped system. Damping is introduced last because, unlike the stiffness and the inertia, it cannot be derived from the geometry, and because the form adopted here is defined in terms of the modes just obtained.

The jacket hanging in air loses energy through several mechanisms at once: internal friction in the steel and Dyneema segments, friction at the sheaves and terminations, contact and damper action within the Block Connection System, and aerodynamic drag on the lattice. None is individually linear in velocity and no single one dominates, so they are not modelled separately but lumped into one equivalent viscous matrix identified from the response of the system itself. The symbol ζ_{air} denotes this combined value rather than a structural contribution alone.

The matrix is constructed in the proportional Rayleigh form $\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}$, which is diagonalised by the same eigenvectors as $\mathbf{M}$ and $\mathbf{K}$. The damped system retains the undamped mode shapes of Table 3.1 and stays modally decoupled, so damping is assigned mode by mode rather than element by element [28]. This is the standard treatment where only a small number of modal damping ratios are known from measurement, which is the situation here: only the sway mode decays cleanly enough to be measured. Releasing the payload in roll, yaw or heave produces transients in which several coordinates ring simultaneously, so no single decaying sinusoid can be isolated. Sway is separated from the next mode by more than a factor of four in period and carries 98.7 % of its own modal kinetic energy, so it decays in isolation.

Nor is this a practical limitation, since sway is the only mode whose damping affects the quantity of interest. Roll and heave lie an order of magnitude from the wave band and respond quasi-statically, where the response is set by stiffness rather than dissipation, and yaw is not driven resonantly but through the off-diagonal stiffness K_{14} at the sway frequency.

Logarithmic-decrement analysis of the sway decay following release from a 2000 kN lateral load gives $\zeta_{\text{air}} = 0.98\,\%$ of critical in the nominal configuration.

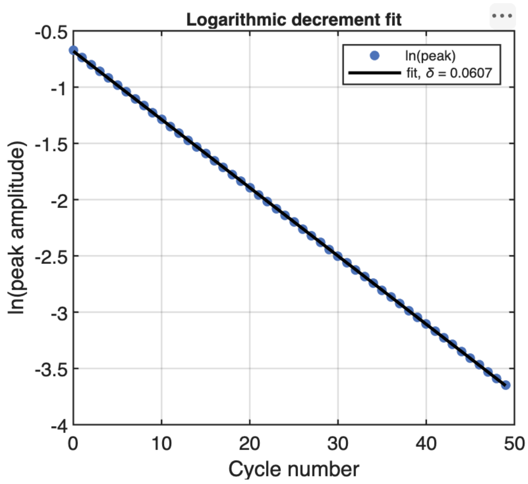

Figure 3.9: Logarithmic-decrement fit for the sway free decay: $\ln$ of successive peak amplitudes against cycle number, for release at 0.54 m. The relation is linear over the full record ($\delta = 0.0607$, $\zeta_{\text{air}} = 0.98\,\%$)

The Rayleigh coefficients follow from

$$\zeta_i = \frac{1}{2}\left(\frac{\alpha}{\omega_i} + \beta\,\omega_i\right), \qquad i = 1, 2, \tag{3.11}$$

which requires two modal ratios and only one is measured. The second anchor is placed at the yaw mode with $\zeta_2 = \zeta_1$, giving constant damping across the two modes that lie within or adjacent to the excitation band and yielding $\alpha = 6.40 \times 10^{-3}$ and $\beta = 9.02 \times 10^{-3}$. The proportional form then fixes the damping of the remaining modes as a consequence rather than a choice: roll and heave receive 2.39 and 2.54 %, since the $\beta\omega$ term grows with frequency. This has no bearing on the results, because both modes lie an order of magnitude above the wave band and respond quasi-statically.

The identified value is carried unchanged through Chapters 5 and 6, so all analyses run on one common plant. At $\omega_n = 0.400\,\mathrm{rad/s}$ it corresponds to a half-power bandwidth of only 0.0078 rad/s. That narrowness is the physically significant result: a resonance this sharp is either excited or it is not, so the response depends strongly on where the excitation sits relative to the mode and comparatively little on how much energy is dissipated within it. This is why detuning is pursued in Chapter 5 before damping augmentation is considered in Chapter 6.

Two qualifications attach to the value. It is identified with the BCS engaged, as the vessel transits, whereas M_{kin} is identified from the bare decay, since the reduced-order model contains no BCS; removing the BCS lowers the identified damping to 0.51 %. The apparent damping also rises with release amplitude, from 0.98 to 1.13 % over the range tested, and with wind speed. Together these place the plausible range at roughly a factor of two in damping and ±3 % in sway frequency. The consequence is a robustness requirement rather than a modelling one, carried forward to Chapter 6. The full free-decay campaign is given in Appendix B.

3.6 Summary of Model Parameters

Table 3.3 summarizes the parameters for the 4-DOF reduced-order model, derived from the physical specifications and system identification. The full numerical mass and stiffness matrices ($\mathbf{M}$, $\mathbf{K}$, $\mathbf{K}_b$), together with the identification of M_{kin} and the sensitivity of the model to the hoist-wire stiffness, are given in Appendix A. The validation of this linear MATLAB model against the non-linear ORCAFLEX reference under various sea states is the subject of Chapter 4.

Symbol	Quantity	Value	Unit
M_{grav}	Gravitational mass (static tensions)	7650	t
M_{kin}	Effective inertial mass (identified)	7719	t
I_{xx}	Roll inertia	4.55×10^6	$\text{t} \cdot \text{m}^2$
I_{zz}	Yaw inertia	3.29×10^6	$\text{t} \cdot \text{m}^2$
k	Wire axial stiffness (per wire)	63 274	kN/m
h_{drop}	Vertical drop, top anchors to padeyes	60.091	m
z_p	Height above jacket CoG to the lifting attachment point	22.41	m
$\overline{L_{\text{eq}}}$	Mean equilibrium wire length	62.515	m
ζ_{air}	Identified sway damping, all contributions	0.98	%
g	Gravitational acceleration	9.81	m/s^2
Δx	Hinge B longitudinal offset from vessel CoG	-215.263	m
Δy	Hinge B transverse offset	1.0	m
Δz	Hinge B vertical offset (above CoG)	99.998	m

Table 3.3: Consolidated summary of parameters for the 4-DOF reduced-order model.

Chapter 4

Model Validation and Baseline Response

This chapter validates the analytical four-degree-of-freedom reduced-order model (ROM) against the fully nonlinear time-domain ORCAFLEX reference. The sequence runs from the least to the most demanding test: first the unforced dynamics, then the steady-state transfer functions, then the response to irregular seas, and finally the response to a measured vessel record. Each stage isolates one property, so that where disagreement appears it can be attributed rather than merely reported.

Every ORCAFLEX run synthesises the wave spectrum from a finite number of components, and the ROM is evaluated at the same number, so every comparison is like-for-like. The effect of that discretisation is quantified in Appendix C.

4.1 Free-Decay Verification

The decay test of Section 3.5 a 2000 kN lateral release at the jacket centre of gravity is repeated on the linear ROM and compared against the ORCAFLEX reference.

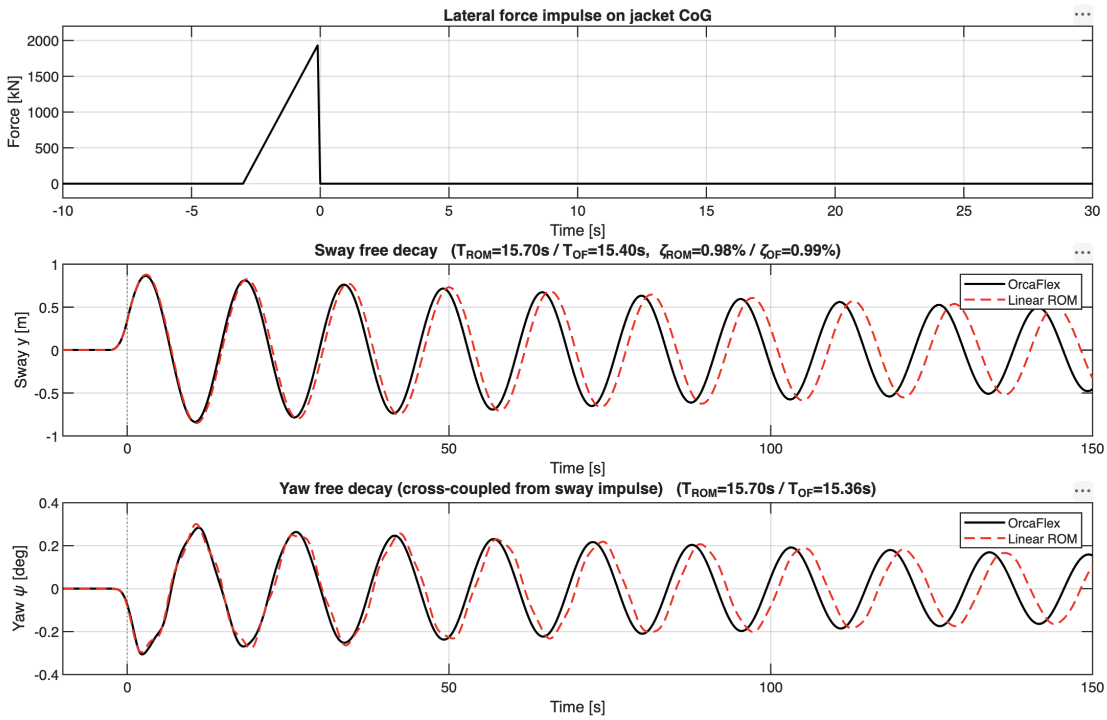

Figure 4.1: Free-decay verification of the sway response and the cross-coupled yaw response. The first 150 s, approximately ten cycles, are shown.

Table 4.1 summarises the comparison.

Table 4.1: Free-decay validation summary.

Metric	Linear ROM	ORCAFLEX	Deviation
Sway period T_1 [s]	15.70	15.40	+1.9%
First peak $\lvert y\rvert$ [m]	0.878	0.863	+1.7%
Damping ratio ζ [%]	0.98	0.99	−1.0%
Yaw period [s]	15.70	15.36	+2.2%

The damping ratio was fitted from this record in Section 3.5, and the inertia was identified from the free-decay period of the bare arrangement. The reference run here has the Block Connection System engaged, as the vessel transits, and the model does not represent it: the BCS restrains the hoist blocks and adds stiffness through that path, shortening the reference period to 15.40 s against the 15.71 s of the bare configuration the model reproduces (Appendix B). The 1.9 % offset accumulates visibly over the record in Figure 4.1 and is the frequency uncertainty the tugger controller is required to tolerate in Chapter 6. The first-peak amplitude entered no identification and agrees to 1.7 %.

The cross-coupled yaw response is the stronger result. The jacket was released in sway alone, yet it yaws , and it does so at the sway period rather than at its own, which identifies the response as driven through the off-diagonal stiffness K_{14} produced by the specific rigging rather than as an independently excited mode. That term is not a tuned parameter. It falls out of the Hessian of the potential at the as-rigged geometry, and the yaw inertia I_{zz} is a mass property of the payload. Nothing in the yaw channel was adjusted to match the reference: the

Rayleigh damping was anchored on the sway mode alone. The predicted yaw amplitude and decay envelope nevertheless track ORCAFLEX throughout, with the same 2 % period offset carried over from the sway channel and the same cause.

4.2 Harmonic Validation: Transfer Functions

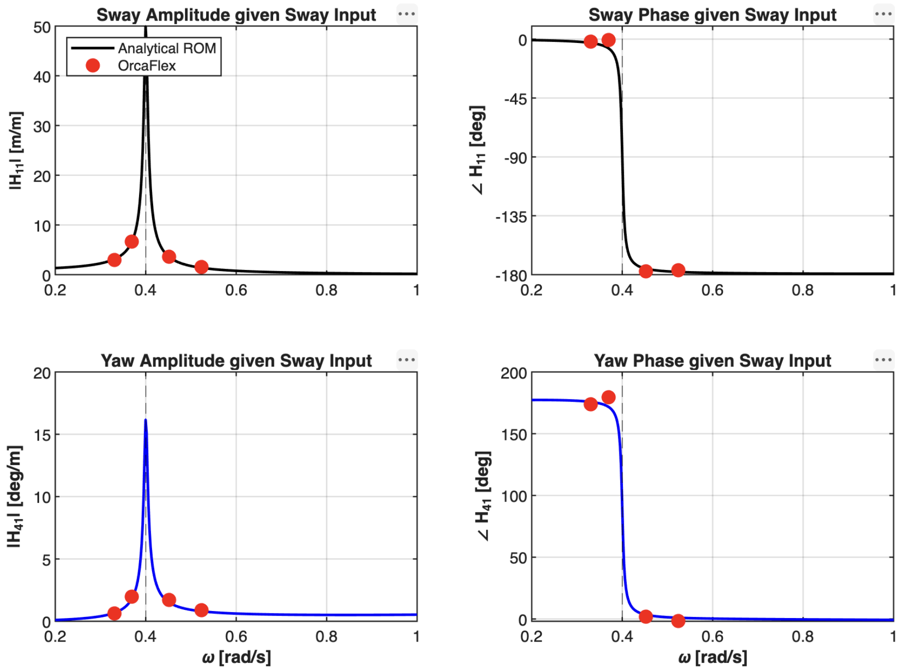

Figure 4.2: Sway and yaw transfer functions of the analytical ROM. Markers give the amplitude and phase identified from the ORCAFLEX steady-state response.

Steady-state harmonic motion of amplitude 0.5 m is prescribed at the hinge at four periods straddling the resonance: 12.0 s and 13.9 s on the sub-resonant flank, 17.0 s and 19.0 s on the super-resonant flank. The empirical points lie on the analytical curve within plotting accuracy and the phase passes through $-90°$ at resonance as it must.

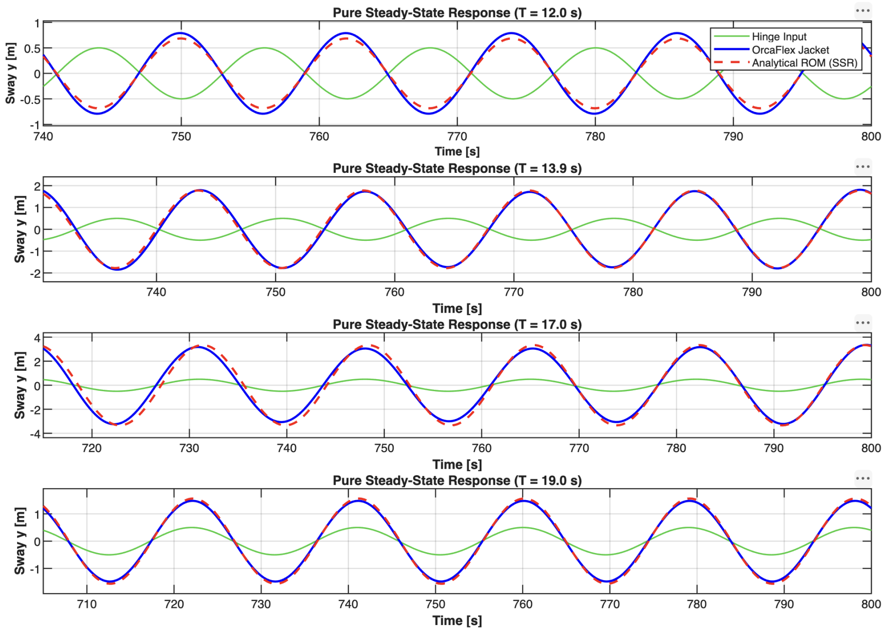

Figure 4.3: Steady-state sway response, analytical ROM against ORCAFLEX, on both flanks of the resonance.

At the exact resonant period the ORCAFLEX amplitude falls below the linear prediction. This is consistent with geometric stiffening: at large swing angles the restoring force grows faster than linearly and the excursion is clamped. The same one-sided over-prediction appears in the irregular-sea results below.

4.3 The Design Envelope and the Critical Heading

The validation set is not arbitrary. It is anchored on the operating envelope the vessel is actually certified for, and on the heading within that envelope which governs.

4.3.1 Beam seas are the limiting heading

Table 4.2 reproduces the current design sea state for North Sea removal projects. The allowable significant wave height is not uniform with heading: it falls from 4.0 m in head and following seas to 2.4 m beam, and the allowable period band narrows with it, from 5.6 s to 8.5 s to 4.3 s to 6.6 s.

Table 4.2: Current design sea state for North Sea removal projects. The allowable wave height is lowest, and the period band narrowest, at 90°.

Table 3 Current design sea state for North Sea removal projects

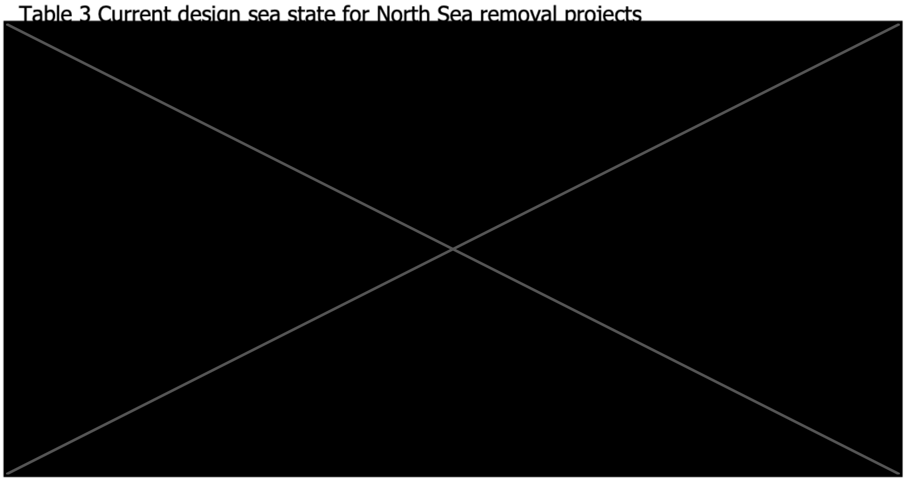

The design table therefore already treats beam seas as the limiting case, restricting H_s at 90° to 60 % of what is allowed head-on. The reason is that the sway mode is driven laterally: by vessel sway, and by roll acting through the 99.998 m lever from the roll axis to the hinge. Both are largest in beam seas.

The reduced-order model is validated at 90° for this reason. It is the heading that limits the operation, and the heading against which the control study is tuned.

4.3.2 The validation set relative to the design cell

The five test sea states are positioned deliberately with respect to the beam sea design transportation state. Two of them are that cell:

- **V1** (H_s = 2.40 m, T_z = 4.30 s) is the lower corner of the 90° beam sea design cell exactly.
- **V2** (H_s = 2.40 m, T_z = 6.60 s) is the upper corner exactly.

The remaining three extend progressively beyond it, so that the model is tested not only where it will be used but also at increasing distance outside:

- **V3** (T_z = 7.50 s) lies beyond the 90° band but within the band allowed at 60°.
- **V4** (T_z = 8.50 s) is at the longest period anywhere in Table 4.2, but carried at the beam-sea wave height rather than the head-sea height that period is certified with.
- **S1** (H_s = 2.70 m, T_z = 9.00 s) lies outside Table 4.2 altogether and is a deliberate stress case.

This ordering is useful when reading the results below: the first two sea states test the model where it will be used, and the last three test how far outside that range it remains reliable.

Figure 4.4 shows the five wave excitation spectra against the sway resonance frequency.

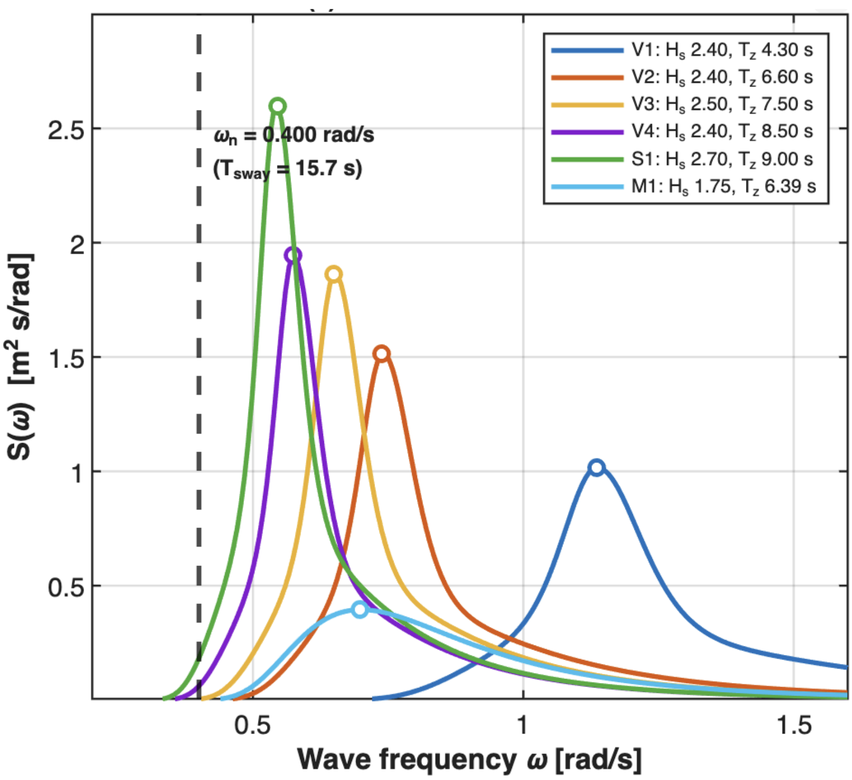

Figure 4.4: Validation sea states against the sway resonance at $\omega_n = 0.400\,\mathrm{rad/s}$.

Table 4.3 lists the characteristic wave periods of the validation sea states relative to the baseline sway natural period (T_{sway}).

Table 4.3: Validation sea states and their characteristic periods relative to the system's sway resonance.

Sea state	H_s [m]	T_z [s]	T_p [s]	T_p/T_{sway}
V1	2.40	4.30	5.53	0.35
V2	2.40	6.60	8.49	0.54
V3	2.50	7.50	9.64	0.61
V4	2.40	8.50	10.93	0.70
S1	2.70	9.00	11.57	0.74

Two primary observations follow. First, all five peak periods lie well below the sway natural period of 15.70 s, meaning the model is nowhere being tested at pure resonance and neither, as later chapters show, is the vessel ever operated there. Second, as T_z and T_p lengthen (progressing toward V4 and S1), the excitation approaches the natural sway period, which drives progressively larger pendular amplitudes and velocities.

Crucially, these larger motions expose the limitations of a linearised model. The discrepancies between the reduced-order model and ORCAFLEX at these higher sea states cannot be attributed solely to the spectral proximity to resonance. Instead, larger physical excursions activate unmodelled nonlinearities that the ROM omits by design. These include the velocity-squared aerodynamic drag, geometric stiffening at large swing angles, and the discrete contact mechanics of the Block Connection System (BCS). Consequently, the longer-period sea states V4 and S1 are identified in advance as the most demanding environments, as they test the boundary where the linear assumption begins to break down.

4.4 Stochastic Validation Against Irregular Seas

Each sea state is synthesised as a JONSWAP spectrum with $\gamma = 3.3$. Both models are driven at the same component count.

4.4.1 Hinge motion

The hinge motion is the model's input. If it did not match, any agreement in relative sway would be two errors cancelling rather than a result. The hinge time trace is therefore extracted from ORCAFLEX and compared against the hinge RMS obtained from the ROM. The ROM reproduces the hinge sway RMS to within 2.7 % at every sea state of meaningful amplitude.

4.4.2 Relative sway

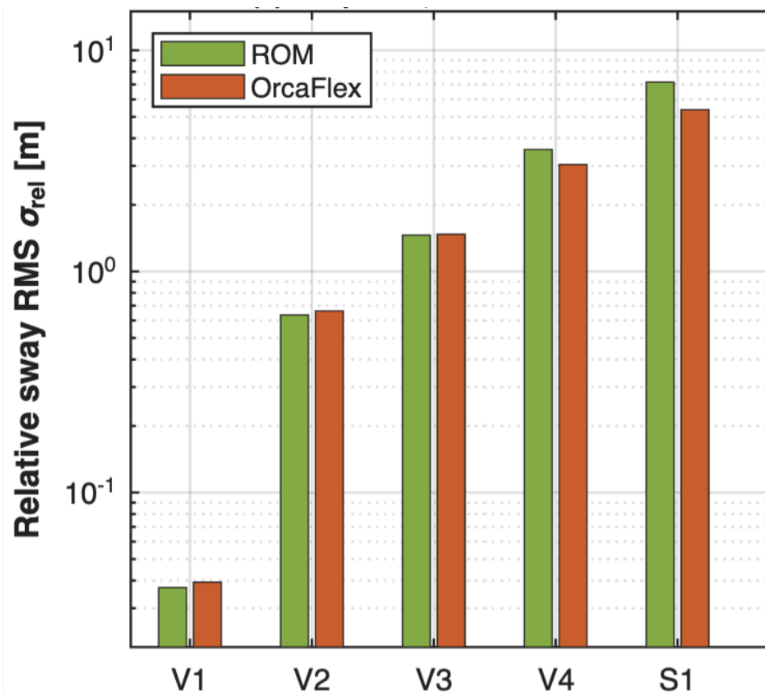

Figure 4.5: Relative sway RMS, reduced-order model against ORCAFLEX.

Table 4.4: Relative sway: RMS against ORCAFLEX, and the three-hour maximum predicted by the model.

Sea state	ROM [m]	ORCAFLEX [m]	Error	Converged σ [m]	3 h max [m]
V1	0.037	0.039	−5.4%	0.060	0.22
V2	0.635	0.663	−4.2%	0.791	**2.94**
V3	1.460	1.472	−0.9%	2.209	8.22
V4	3.560	3.047	+16.8%	6.144	22.86*
S1	7.178	5.383	+33.4%	10.042	37.36*

* At V4 and S1 the linearised model over-predicts and the excursions exceed the range in which it is valid. These figures indicate the trend, not the physical response.

The first three columns compare the two models at the same number of wave components, which is what makes the comparison like-for-like. The fourth is the model's own converged result, evaluated on the continuous spectrum, and it is the one carried forward: the last column is the largest excursion expected in three hours of operation, $y_{\max} = 3.72\,\sigma$, the Rayleigh most probable maximum for 1000 response cycles. This is the quantity the clearance budget limits, and it is what the hoist and tugger chapters are judged on.

The result splits cleanly along the boundary of the design envelope.

- **Inside the 90° design cell** V1 and V2, its two corners the model agrees to within 5.4 %, and does so on both sides of zero.
- **At V3**, just beyond the 90° band but still inside the band allowed at 60°, agreement is 0.9 %.
- **At V4 and S1**, at and beyond the longest period in Table 4.2, the model over-predicts by 16.8 % and 33.4 %.

The agreement is good inside the design envelope and degrades as the excitation is pushed past it toward the sway resonance. V4 and S1 are also the two sea states carrying by far the most response energy near the resonance, so the disagreement follows the resonant content across the set.

4.5 Time-Domain Comparison at Identical Excitation

In the comparison above each model computes its own hinge motion. Both begin from the same wave spectrum and the same vessel response amplitude operators, but each discretises that spectrum into its own set of wave components, so the hinge motions they produce are not identical only statistically the same.

A more direct test is to hand the reduced-order model the hinge motion that ORCAFLEX computed, instead of letting it construct its own. Both models then respond to exactly the same hinge displacement at every time step, and any difference that remains comes from the plant.

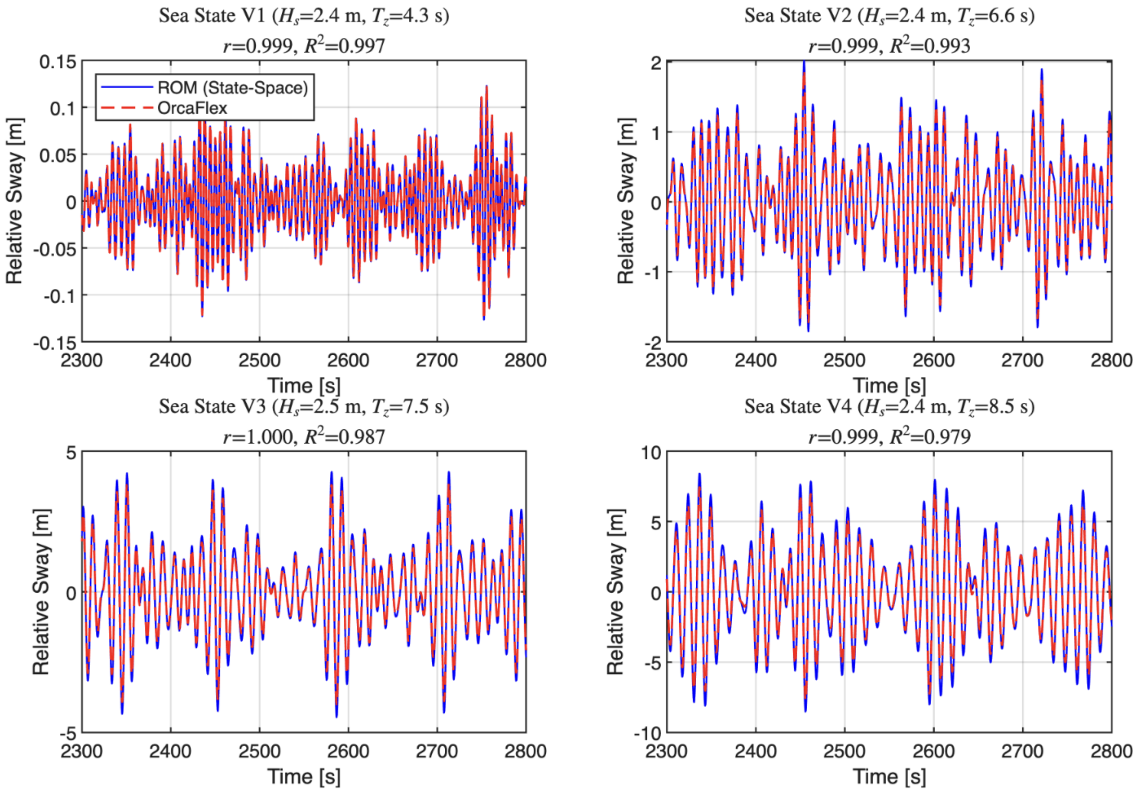

Figure 4.6: Relative sway, reduced-order model against ORCAFLEX, both driven by the same hinge record.

Table 4.5: Time-domain validation, ORCAFLEX hinge record as input to both models.

Sea state	ROM [m]	ORCAFLEX [m]	RMS error	r	R^2
V1	0.040	0.039	+0.6%	0.9986	0.9972
V2	0.715	0.663	+7.9%	0.9995	0.9927
V3	1.637	1.472	+11.2%	0.9996	0.9865
V4	3.451	3.047	+13.3%	0.9986	0.9792
S1	6.117	5.383	+13.6%	0.9948	0.9697

The correlation coefficient exceeds 0.994 and the coefficient of determination exceeds 0.969 at every sea state, including the two where the spectral comparison agreed least. The model reproduces the shape of the signal, not only its statistics.

4.6 Validation Against Measured Vessel Motion

The tests so far compare one model against another, all of them driven by idealised JONSWAP spectra. The strongest available check replaces the synthesised excitation with the measured vessel record from an actual transit: 5389.8 s of roll, pitch and yaw, high-pass filtered at 0.02 Hz to remove static heading drift, transferred to the hinge through the measured lever arms.

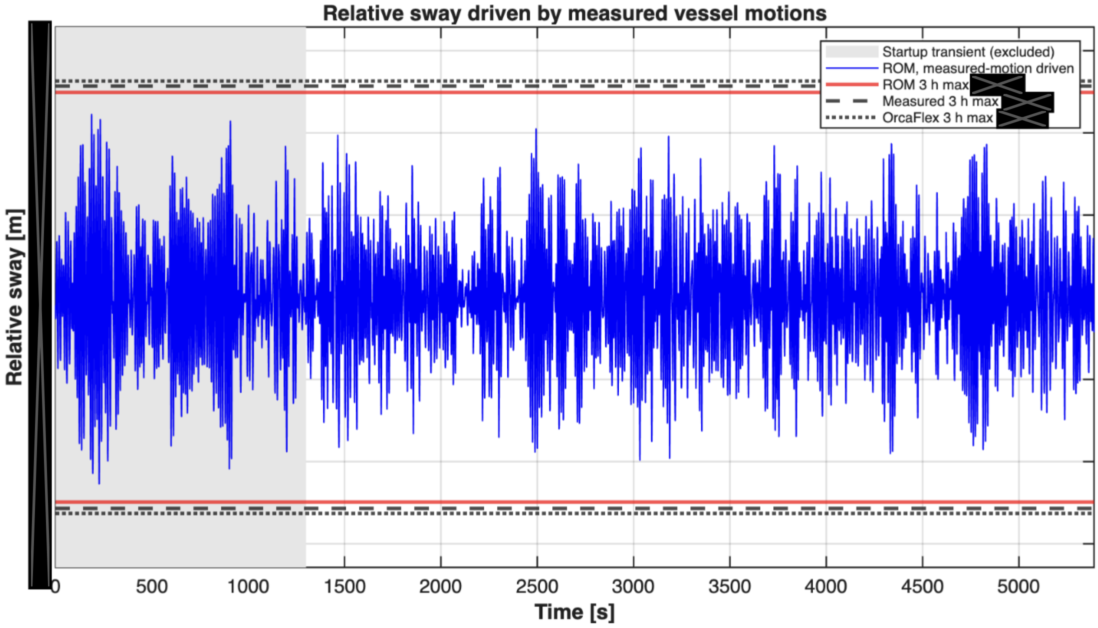

Figure 4.7: Relative sway predicted from the measured vessel motion record.

The simulation begins from rest, and with $\zeta = 0.98\,\%$ at 15.70 s the sway mode settles with a time constant of 255 s. The first 1300 s, approximately five time constants, are therefore excluded from the statistics.

Over the remaining record the standard deviation of relative sway is [redacted] giving a three-hour most probable maximum o [redacted]

Table 4.6: Three-hour most probable maximum relative sway, measured transit record.

Source	MPM [m]	ROM deviation
Reduced-order model	[redacted]	—
Measured	[redacted]	-3.0%
ORCAFLEX	[redacted]	-5.2%

The model predicts the three-hour maximum of a real operation to within 3 %. This is the only comparison in this work against measured data rather than against another model, and it is the result that most directly supports the use of the ROM for operability assessment.

4.7 Suitability of the ROM for Design and Control

The validation establishes four things.

The unforced dynamics are correct. Released and left to swing, the model reproduces the reference decay envelope, with the first peak agreeing to 1.7 %.

The forced response is correct in waveform, not merely in statistics. Correlation exceeds 0.994 and R^2 exceeds 0.969 at every sea state when both models are driven by the same excitation.

The disagreement lies outside the design envelope, is bounded, and is one-signed. Inside the beam sea design envelope, at both its corners, agreement is within 5.4 %. Beyond it the model over-predicts, and no single cause accounts for the whole of it. Three contribute. The linear plant carries a single constant modal damping ratio, whereas the free-decay evidence of Appendix B shows the dissipation rising with amplitude and with engagement of the Block Connection System; the effective damping implied by the ORCAFLEX response at the longest periods is higher than any value identified in free decay. Geometric stiffening at large swing angles, which the linearised restoring term does not carry, adds to this. The finite wave-component count contributes a further resolution effect rather than a modelling error (Appendix C). All three act in the same direction, so the ROM over-predicts and acts as a conservative upper bound.

The prediction holds against measured data. Driven by a real transit record, the model reproduces the three-hour maximum to within 3 %.

These properties suit the two uses the model is put to. For evaluating rigging geometries, a conservative upper bound is what is wanted, and the comparison between configurations is in any case a ratio in which the common bias largely cancels. For control synthesis, the relevant regime is the one in which the controller keeps the payload small displacements, well inside the linearised range where the model is at its most accurate.

Chapter 5

Passive Hoist Optimisation

The validated reduced-order model (ROM) provides a computationally efficient means of assessing how changes in hoist geometry influence the suspended jacket dynamics. This enables a passive optimisation study in which the rigging arrangement is adjusted to detune the dominant sway mode away from the design transportation state excitations, thereby reducing resonance amplification without requiring active actuation. The optimisation is conducted on the linear ROM established in Chapter 3 and validated in Chapter 4, with the nonlinear ORCAFLEX model retained for final verification .

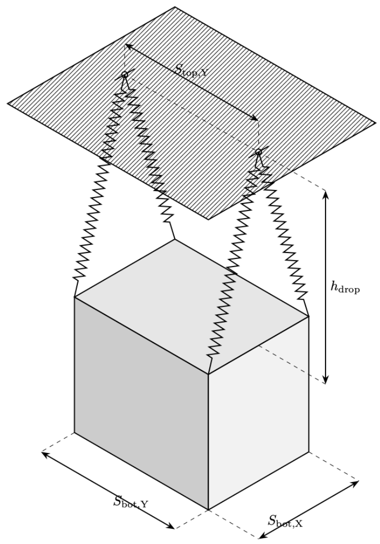

Figure 5.1: 3D Isometric View

The primary objective of the passive study is to minimise the resonant response of the suspended jacket by shifting the sway natural period away from the design wave periods encountered during transit. In practice, this is achieved by modifying the hoist geometry, most notably the top spread and the drop height, while respecting operational and geometric constraints imposed by the Jacket Lift System (JLS). The design target is not to eliminate motion entirely, but to reduce the transfer of wave-induced energy into the payload by lowering the overlap between the resonances and the wave excitations.

5.1 Design variables and constraints

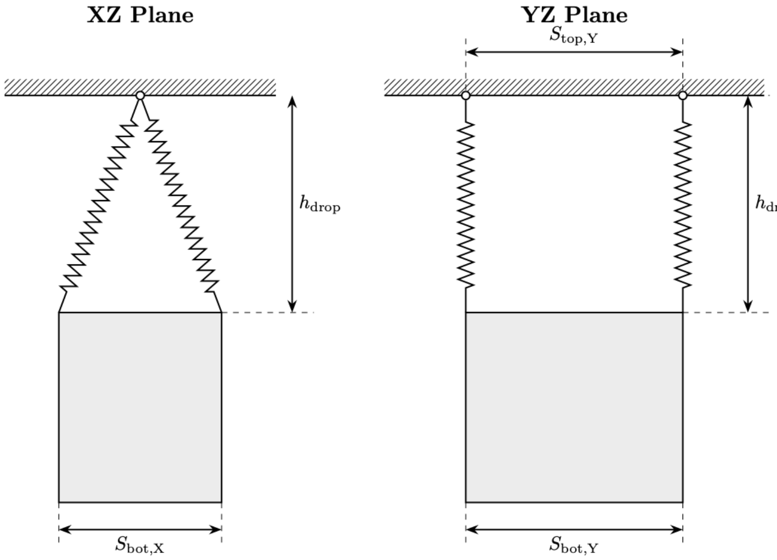

Figure 5.2: Hoist Parameters

The hoist parameters that most strongly affect the pendular restoring stiffness are:

- Top spread Y, $S_{\text{top,Y}}$.
- Bottom spread Y, $S_{\text{bot,Y}}$.
- Bottom Spread X,$S_{\text{bot,X}}$.
- Drop height, h_{drop}.

The feasible range of each variable is set by the JLS specifications, Jacket footprint and by rigging and clearance limits, summarised in Table 5.1. The drop-height ceiling follows from the hinge elevation of $119.46\,\text{m}$, the $48\,\text{m}$ from padeye to jacket tip and a required splash clearance of $5.0\,\text{m}$, leaving only $6.4\,\text{m}$ of pay-out room relative to the baseline.

Table 5.1: Design variable bounds and rigging constraints.

Quantity	Bound	Origin
Top spread, $S_{\text{top,Y}}$	10.5–$36.0\,\text{m}$	JLS geometry
Bottom spread, $S_{\text{bot,Y}}$	30–$47\,\text{m}$	jacket padeye layout
Bottom spread, $S_{\text{bot,X}}$	10–$34.2\,\text{m}$	jacket padeye layout
Drop height, h_{drop}	50.0–$66.5\,\text{m}$	splash clearance
Side (athwartships) lead angle	≤ 10°	rigging
Fore–aft lead angle	≤ 50°	rigging

5.2 Passive detuning and Sensitivity

For small oscillations, the suspended jacket behaves approximately as a pendular system whose sway stiffness is governed by the effective gravitational restoring action of the hoist arrangement. Increasing the drop height reduces the restoring stiffness and increases the sway period, while reducing the top spread alters the line geometry and modifies the coupling between sway and yaw. The design idea is therefore to move the dominant sway resonance away from the main wave-energy peak so that the operational sea-state band excites the system less efficiently .

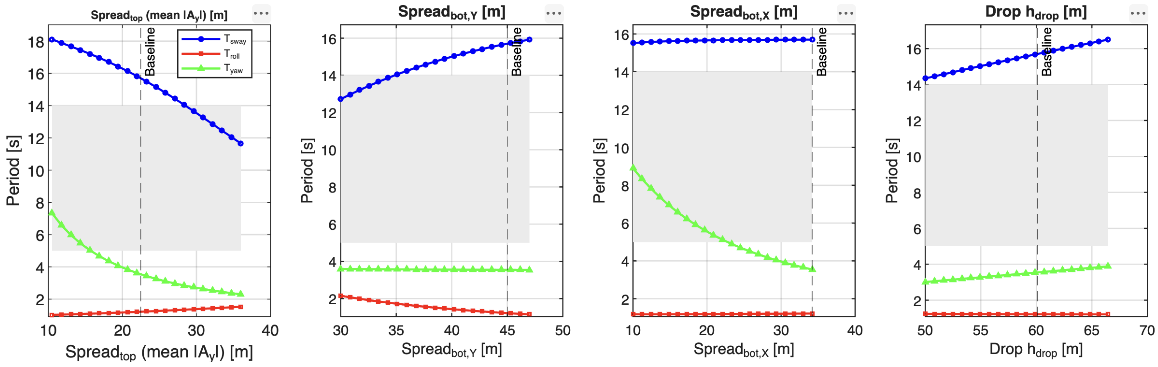

Figure 5.3: Sensitivity of the coupled natural periods to each hoist design variable, computed with the 4-DOF ROM by varying one variable at a time about the baseline arrangement (dashed).

Figure 5.3 sweeps each variable in turn about the baseline arrangement. Three observations follow. The top spread is by far the strongest lever on the sway period, which rises from 11.5 s to over 18 s as the spread is narrowed across its feasible range. The drop height acts in the same direction but not as strong, and is in any case bounded to 6.4 m of pay-out by the splash-clearance requirement. The bottom spread in X leaves the sway period essentially unchanged while moving the yaw period from 9 s to 3.5 s.

The same figure shows the cost of the detuning mechanism. Narrowing the top spread raises the sway period out of the wave band, but it also raises the yaw period from below 4 s towards 7 s, in the direction of the band rather than away from it. The two effects are geometrically inseparable, and the consequence is examined in Section 5.7.

Furthermore, evaluating the operational margins relative to the baseline arrangement in Figure 5.3 reveals the practical limits of passive detuning. While both the top spread and drop height theoretically allow for an upward shift in T_{sway}, the drop height is heavily restricted by the aforementioned splash-clearance constraint. Consequently, narrowing the top spread provides the largest and most viable constraint space to effectively detune the system away from the wave-energy peak, despite the geometric penalty imposed on the yaw period.

5.3 Optimisation formulation

The passive hoist problem is formulated to minimise the variance of the payload sway motion. Because the linearised system response is narrow-banded, the Most Probable Maximum (MPM) of the sway excursion is directly proportional to the standard deviation (RMS) of the response through a statistical peak factor. Therefore, minimising the root-mean-square of the sway motion inherently minimises the extreme peak excursions without requiring arbitrary weighting parameters. The optimisation is defined as:

$$
\begin{aligned}
\underset{\mathbf{x}}{\text{minimise}} \quad & J(\mathbf{x}) = \mathrm{RMS}_{\text{sway}}(\mathbf{x}) \\
\text{subject to} \quad & T_{\text{sway}}(\mathbf{x}) \notin B_{\text{wave}}, \\
& \mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max}, \\
& \text{geometric and clearance constraints.}
\end{aligned}
\tag{5.1}
$$

where $\mathbf{x}$ denotes the vector of hoist design variables (top spread, bottom spread, and drop height), T_{sway} is the dominant sway period, and B_{wave} is the design wave-period band. By directly minimising the RMS, the geometric restoring stiffness is systematically tuned to shift the system's dominant resonances away from the expected wave energy.

The objective is written in terms of sway alone. This is justified by the behaviour of the baseline plant, in which yaw is not an independently excited mode: the reduced-order model places the yaw natural period at 3.54 s, far below the wave band, and the relative yaw transfer function peaks at the sway natural period of 15.71 s rather than at its own. The ORCAFLEX records confirm this independently. Across every baseline sea state the relative yaw and relative sway time series are correlated at $|r| \geq 0.995$, with a yaw-to-sway gain of 0.27–0.33°/m that is constant across sea states and across wave period. Yaw is therefore a kinematic consequence of sway, and minimising sway RMS minimises both. Whether this remains true once the geometry is altered is not assumed here; it is later checked in Section 5.7 where we analyse the system in optimised geometry in Orcaflex.

5.4 Screening of the design space

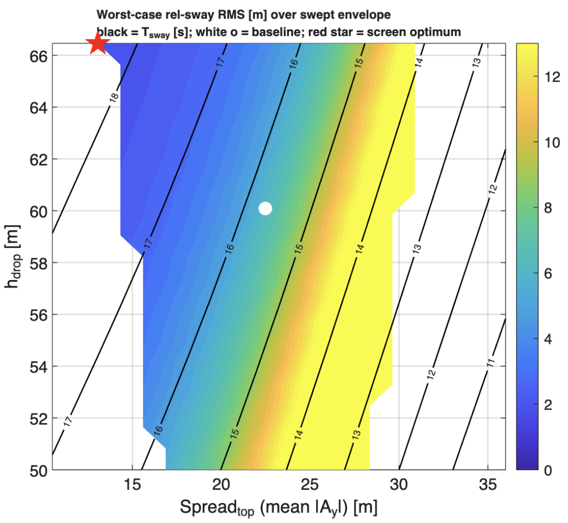

Figure 5.4: Screening surface for the passive hoist optimisation: worst-case relative sway RMS over the swept sea-state envelope as a function of top spread and drop height, with sway natural period contours overlaid. White marker: baseline. Red marker: screening optimum at the constraint boundary. White areas are geometrically infeasible due to the constraints.

The screening envelope used throughout this chapter sweeps T_z from 5.5 to 8.5 s at a constant $H_s = 2.5\,\mathrm{m}$. This is deliberately broader and slightly more severe than the beam-sea design cell of the transportation memo ($H_s = 2.4\,\mathrm{m}$, $T_z \leq 6.6\,\mathrm{s}$) against which Chapter 6 evaluates the tugger: the extra period range covers the region where the baseline sway resonance is approached, and the optimisation is therefore graded on a more demanding basis than the design cell alone would give.

Figure 5.4 maps the worst-case relative sway RMS across the swept sea-state envelope over the feasible range of top spread and drop height, with contours of the corresponding sway natural period overlaid. Two features are apparent. First, the colour gradient runs almost parallel to the T_{sway} isolines: the objective is governed by the natural period alone, and the two design variables influence the response only through their combined effect on the pendular restoring stiffness. This confirms that the optimisation is a detuning problem rather than a two-parameter trade-off, and it is the reason a single scalar objective is sufficient. Second, the top spread is the more effective of the two variables, since the isolines are crossed far more rapidly by reducing spread than by increasing drop height. This is reflected in the optimised configuration, which reduces the top spread by 42% while increasing the drop height by only 11%.

Within the feasible domain the objective decreases monotonically with increasing T_{sway}, so no interior minimum exists and the optimum is located at the corner of the feasible region where the geometric and clearance constraints of the JLS become active. The baseline arrangement (white marker, $T_{\mathrm{sway}} = 15.70\,\mathrm{s}$) lies in the region of high response, while the screening optimum (red marker, $T_{\mathrm{sway}} = 18.39\,\mathrm{s}$) lies at the constraint boundary. The optimum is therefore constraint-limited rather than physics-limited: further detuning would continue to reduce the response, and the achievable improvement is set by how far the rigging geometry is permitted to move.

5.5 Optimised configuration

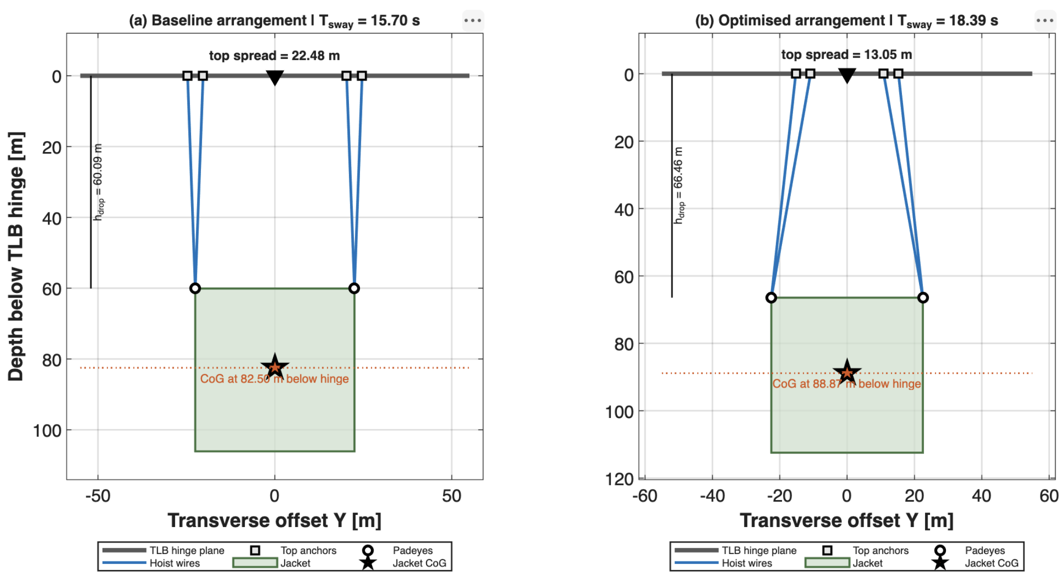

Figure 5.5: Transverse (Y–Z) projection of the baseline (a) and optimised (b) hoist arrangements, viewed along the vessel axis.

The optimisation identifies a configuration that increases the sway natural period from the baseline value of 15.70 s to 18.39 s. This shift moves the dominant resonance away from the main swell band and yields a substantial improvement in operability across the analysed design sea states. The corresponding geometry uses a reduced top spread and increased drop height relative to the baseline arrangement, consistent with the physical detuning mechanism captured by the ROM .

Table 5.2: Baseline and optimised hoist configuration from ROM Model.

Parameter	Baseline	Optimised
Top spread, S_{top} [m]	22.48	13.05
Drop height, h_{drop} [m]	60.09	66.46
Sway period, T_{sway} [s]	15.70	18.39

5.6 Effect on response and operability

Figure 5.6 shows the reduction in relative sway RMS across the wave-period range at constant H_s. The benefit is largest near the baseline resonance, where the original geometry exhibits the highest dynamic amplification, and remains positive across the design band. Because the optimised plant is detuned, the response curve is flatter and the peak amplitude is lower across the operational sea states of interest.

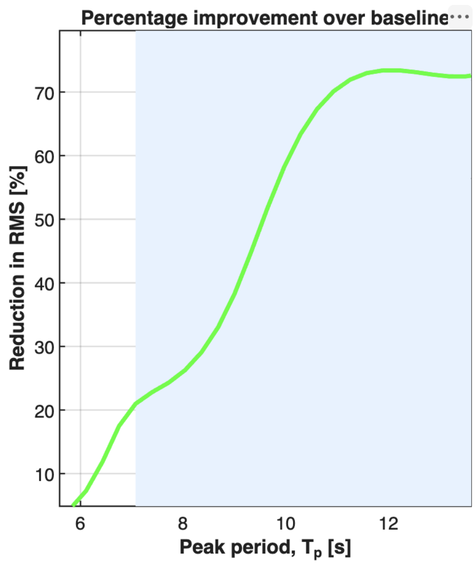

Figure 5.6: Relative sway RMS percentage reduction for the baseline and optimised arrangements across the wave-period range at constant H_s.

Figure 5.7 shows the mechanism directly. Detuning moves the sway resonance from 15.70 s to 18.39 s, clear of the design excitation band, and the peak amplitude falls by roughly a factor of three.

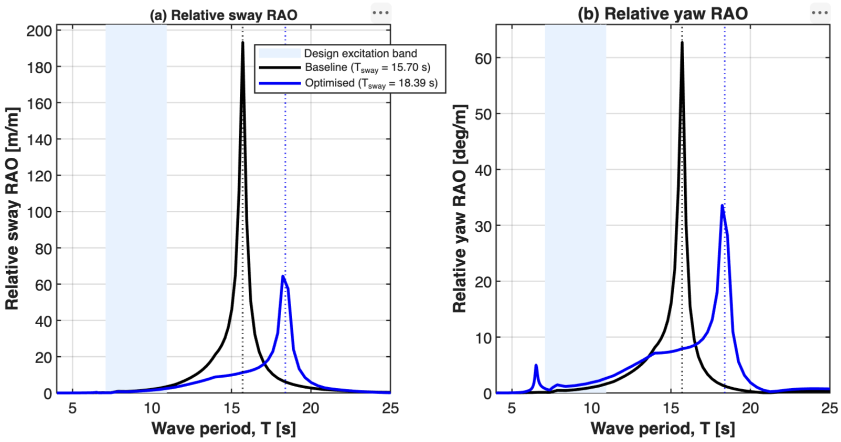

Figure 5.7: Relative sway (a) and relative yaw (b) transfer functions for the baseline and optimised arrangements. Dotted lines mark T_{sway}.

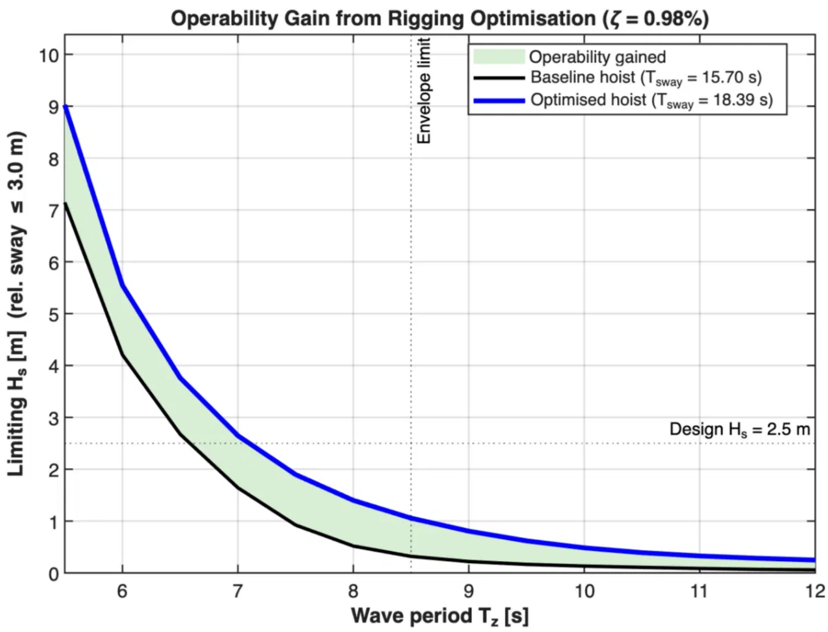

Figure 5.8: Operability gain from the rigging optimisation, expressed as the limiting significant wave height for a relative sway criterion of $3.0\,\text{m}$ and a modal damping ratio of $\zeta = 0.98\%$.

Figure 5.8 presents the resulting operability gain, expressed as the limiting significant wave height at which the relative sway criterion of $3.0\,\text{m}$ is reached, as a function of the zero-crossing period. The optimised arrangement raises the limiting H_s across the entire period range. The gain is largest in relative terms at long periods, where the baseline response is dominated by proximity to its resonance: at the upper end of the swept band, $T_z = 8.5\,\text{s}$, the limiting H_s increases from approximately $0.35\,\text{m}$ to approximately $1.05\,\text{m}$, a factor of three. In terms of workable period range, the design condition of $H_s = 2.5\,\text{m}$ is sustained up to $T_z \approx 6.6\,\text{s}$ for the baseline arrangement and $T_z \approx 7.1\,\text{s}$ for the optimised arrangement.

The limitation of the passive approach is equally apparent. Although the optimisation triples the limiting wave height at $T_z = 8.5\,\text{s}$, the resulting value remains well below the design condition of $H_s = 2.5\,\mathrm{m}$. Detuning the rigging geometry therefore provides a substantial but partial improvement: it reduces the response for a given sea state without extending operability to the full design envelope at the longer periods. Closing the remaining gap requires the introduction of additional damping rather than a further change in restoring stiffness, which motivates the active tugger-line control study of Chapter 6.

A caveat applies to the interpretation of these gains. The optimised arrangement is obtained by altering the hoist geometry, and any change in sling inclination redistributes the static tension between the hoist wires and modifies the load introduction at the jacket lift points. The reduced motions reported above lower the inertial contribution to the total load, while the modified geometry changes the static contribution, so the net effect on structural utilisation is not determined by the motion results alone. The 4-DOF ROM represents the jacket as a rigid body and therefore cannot resolve internal member forces or lift-point reactions by construction; the ORCAFLEX verification model shares this limitation for the payload itself. Confirming that the optimised configuration remains within the allowable sling tensions and lift-point capacities requires a dedicated structural assessment of the jacket and its rigging, which falls outside the scope of the present dynamic study. The optimised geometry is therefore presented as a dynamically favourable configuration, subject to structural verification before it could be adopted operationally.

5.7 Verification in ORCAFLEX

To confirm that the performance gains predicted by the linear reduced-order model (ROM) translate to the physical system, the optimised geometry is re-evaluated in the full nonlinear, time-domain ORCAFLEX model.

5.7.1 Verification of Modal Detuning

Before evaluating the time-domain response, the fundamental mechanism of the optimisation ie the shift in natural periods is verified. A modal analysis of the modified geometry was performed in the full 3D ORCAFLEX model to confirm that the geometric detuning predicted by the 4-DOF ROM physically manifests.

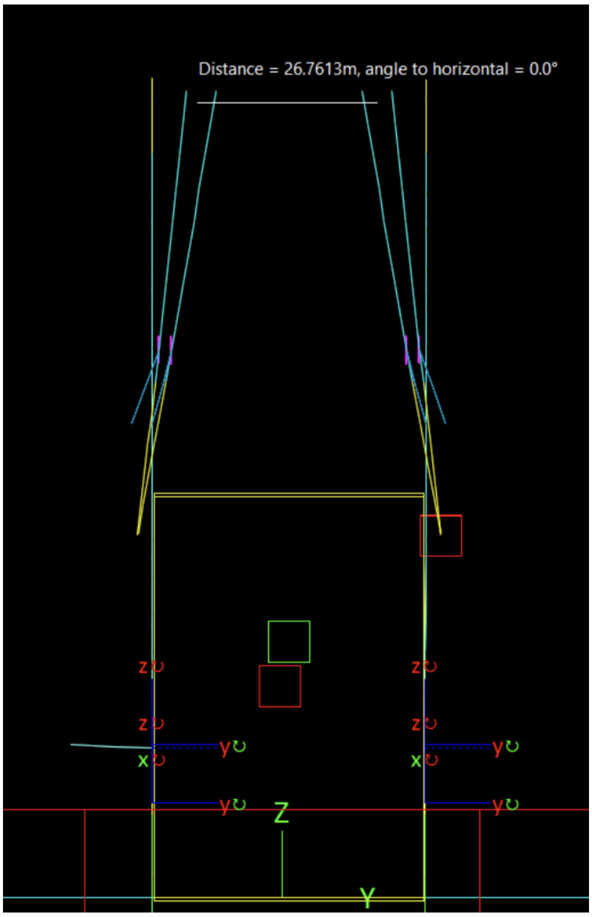

Figure 5.9: Implementation of the optimised hoist geometry in the ORCAFLEX model

The ROM predicts a detuned sway period of $T_{\mathrm{sway}} = 18.39\,\mathrm{s}$. The ORCAFLEX modal analysis confirms this, identifying the dominant sway mode at $18.56\,\mathrm{s}$ (Mode 2). This excellent agreement (less than 1% deviation) demonstrates that the ROM accurately captures the gravitational restoring stiffness of the new multi-wire spatial arrangement. Note that Mode 1 in the 3D model ($20.33\,\mathrm{s}$) corresponds to out-of-plane longitudinal surge, which is inherently excluded from the beam-sea 4-DOF framework.

Furthermore, the ROM predicts that narrowing the top spread softens the yaw stiffness, increasing the yaw period to $T_{\mathrm{yaw}} = 6.52\,\mathrm{s}$. The ORCAFLEX analysis validates this yaw softening trend, placing the new yaw-dominated mode at $7.64\,\mathrm{s}$ (Mode 3). While the full 3D model exhibits slightly more yaw stiffness than the simplified 4-DOF representation, both models unequivocally confirm the same physical consequence: the yaw mode is pushed upward and closer to the design wave band.

5.7.2 Verification of Sway Reduction

Figure 5.10 compares the relative sway response and the percentage reduction achieved by the optimised rigging against the baseline.

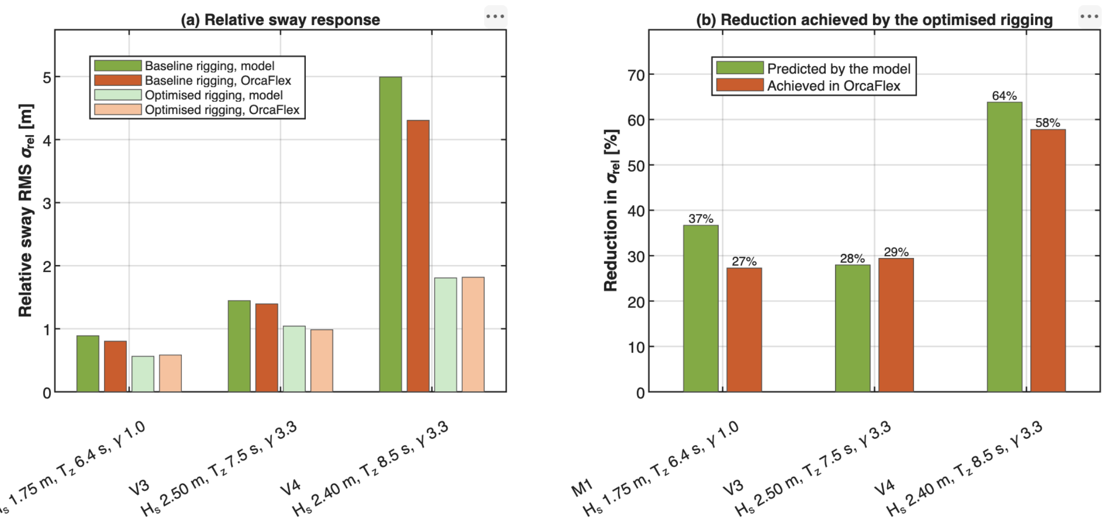

Figure 5.10: Comparison of the relative sway RMS (left) and the percentage reduction (right) predicted by the 4-DOF model and achieved in the fully nonlinear ORCAFLEX model.

The reduction achieved by the optimised rigging grows with wave period across the two design states: 29% at V3 (T_z = 7.5 s) and 58% at V4 (T_z = 8.5 s). This is the expected signature of detuning, since the benefit is largest where the baseline sway resonance sat closest to the excitation. The ROM reproduces this trend and matches the ORCAFLEX reduction to within 1.4 pp at V3, while overpredicting it by 6 pp at V4. The measured sea state M1 (H_s = 1.75 m, T_z = 6.4 s, γ = 1.0) provides an independent check against a milder and broader-banded spectrum rather than a third point on the same trend; there the ROM predicts 37$\%$ against 27$\%$ achieved.

In absolute terms the ROM overpredicts the baseline sway at V4. This bias acts on both geometries and therefore largely cancels in the ratio, but not entirely, and the residual is one-directional: the linear model consistently flatters the optimisation. The ROM is accordingly used as a screening tool for ranking hoist configurations, with its predicted reductions treated as upper bounds pending time-domain confirmation. The V3 comparison is made with the BCS disengaged and on a coarser wave discretisation than the other two.

5.7.3 Other DoF motions

Narrowing the top spread softens the yaw restoring stiffness, raising the yaw natural period from 3.54 s to 6.52 s in the ROM and to 7.64 s in the ORCAFLEX modal analysis, from well below the design wave band to its lower edge. The slaving assumed in Section 5.3 is therefore weakened, but not removed: the yaw transfer function of the optimised arrangement still peaks at around its sway period of 18.23 s rather than at 6.52 s.

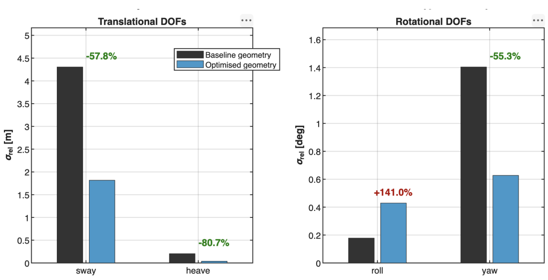

Figure 5.11: Relative response in all four modelled degrees of freedom at the V4 sea state, baseline against optimised geometry. OrcaFlex runs

Figure 5.11 widens the comparison from sway to all four modelled degrees of freedom at the fully matched V4 run. Relative sway falls from 4.305 m to 1.817 m and relative yaw from 1.404° to 0.627°, both by more than half, so the two motions that carry the clearance argument move together as the slaving assumption requires. Relative heave and relative roll are an order of magnitude smaller in both geometries heave 0.203 m to 0.039 m, roll 0.178° to 0.429° and the percentages attached to them should be read with that in mind. A ratio of two small quantities is sensitive both to the realisation and to the residual drift in the record, and neither response approaches an operational limit in either arrangement.

Roll is the only degree of freedom that moves in the unfavorable direction. It is governed by the axial stiffness of the hoist wires rather than by gravitational restoring, so T_{roll} is 1.21 s at the baseline and 1.06 s in the optimised arrangement: an order of magnitude below the wave band in both, and stiffened rather than softened by the narrower top spread. The rise in the relative quantity is therefore not a resonance effect, and it is not traced to a single mechanism within the scope of this study. On this payload the consequence is immaterial, the response remaining under half a degree and unconstrained by any criterion applied here. What generalises is not the magnitude but the direction: an objective written on sway alone offers no protection to the rotational modes, and how much that costs depends on the padeye layout and mass distribution of the payload being rigged. That is the argument for carrying them as explicit constraints in any generalised form of the method, developed in Section 5.8.

5.8 Conclusion

The passive optimisation delivers a practical geometry that improves operability with no additional hardware, no control system and no crew intervention: it is a matter of how the jacket is rigged before the transit begins. It also defines the best passive plant against which an active device can be judged, since any damping system fitted to this payload must be shown to earn its place relative to a rigging change that costs nothing to operate. Whether it does so is the subject of Chapter 6, which sizes a damping tugger on the baseline plant and then places both mitigations on the same axes.

The optimisation presented here minimises relative sway alone, on the grounds that yaw follows sway at a fixed gain in the baseline plant and is not independently excited. The verification bears that out for the arrangement as sailed: with the block control system engaged, sway and yaw fall together at every sea state examined, and the yaw carried per metre of sway is essentially unchanged by the geometry.

The assumption is nonetheless doing more work than it appears. Narrowing the top spread softens the yaw mode from well below the wave band to its lower edge, and the consequence becomes visible once the restraint hardware is removed: with the BCS disengaged the optimised geometry develops a distinct relative-yaw peak at the new rotational period and a higher relative yaw than the baseline. The reduced-order model carries no representation of that hardware, so its rotational predictions describe the unequipped plant. What keeps the penalty from appearing in service is a subsystem the optimisation never modelled, and the favourable outcome is therefore verified against the nonlinear model rather than derived from the objective.

That is an argument for a broader formulation, not a broader claim. A single-degree-of-freedom objective is adequate here because this payload happens to couple favourably and because the equipped vessel absorbs what the objective ignores. None of that is guaranteed for another payload. A different jacket, a different padeye layout, or a vessel without the same restraint hardware could behave differently, and the present formulation would give no warning. Including the rotational motions as constraints and later verifying with the Orcaflex model, would make the outcome something the method delivers rather than something this particular system happens to allow.

Towards a route-specific rigging assessment tool

Nothing in the method is specific to this jacket. The design variables, the constraints and the objective are all built from data available before a transit: the jacket's mass, inertia and padeye layout, the vessel's response operators, and a description of the sea. Only the last was fixed here, to the conservative design envelope of section 4.3. Replace it with the wave scatter of a planned voyage and the same calculation returns the best rigging for that voyage, together with the fraction of the scatter in which the clearance criterion is met and quickly, since the search runs on the reduced-order model and only the selected geometry needs a nonlinear check. Two additions would make it usable in that form: sling tension and lift-point capacity as constraints rather than rigging geometry alone, and the rotational motions carried explicitly instead of assumed to follow sway.

Chapter 6

Tugger Control

The previous chapter demonstrated the reduction of jacket sway through passive adjustments to the rigging geometry. This chapter investigates whether that sway motion can instead be reduced actively by extracting energy through tugger lines. Finally, the passive and active methods are compared and combined to evaluate their relative effectiveness.

The three control modes available to a tugger—fixed lines, constant tension, and active damping—were reviewed in Section 2.6.1. Section 6.2 selects the appropriate mode for the present system on quantitative grounds. The remainder of the chapter formulates the selected control law, scales it into three configurations based on the JLS architecture, and evaluates their performance across the design transportation states and upon widening these. The objective is to establish the dynamic conditions under which a damping tugger is worth fitting to a payload of this scale, rather than to specify a finalized winch design for a single jacket.

6.1 Control objective and success criterion

The success criterion is inherited from the uncontrolled study: the most probable maximum (MPM) of the relative sway over a three-hour exposure must remain within the clearance budget of 3.0 m, a limit set by the operator. This is a constraint, not an objective to be minimised: the question is whether a sea state can be worked, not how much motion can be removed.

One property of the system frames everything that follows. The sway period of the jacket is 15.70 s. A sea state whose peak period matches it has a zero-crossing period of about 12.2 s (JONSWAP, $\gamma = 3.3$ [29]), whereas the current beam-sea design transportation envelope ends at $T_z = 6.6\,\mathrm{s}$. The jacket is therefore excited well away from its own sway period across the whole design box.

The response there has two parts. Most of it follows the motion of the suspension point directly, at the wave period. The remainder is oscillation about that point at the sway period, fed by the long-period tail of the spectrum and lightly damped at 0.98 %; it arrives in bursts, and it is what sets the three-hour maximum rather than the typical amplitude. A tugger applies force at every frequency present in the motion, but only the second part is damping-controlled: the first is set by the motion of the suspension point and is largely insensitive to how much damping is added. As the wave period lengthens toward 15.70 s a larger share of the spectrum falls near the sway period and that second part grows. The effectiveness of a damping tugger should therefore be smallest at the upper period limit of the present beam-sea envelope and should rise steadily beyond it.

6.2 Tugger types

6.2.1 Fixed tugger

Fixed lines are excluded on two independent grounds, either of which is sufficient on its own.

The first is the load. The argument is due to Meskers and Dijk [19] and is reproduced here for the present payload. A line held at a fixed length forces the jacket to follow the horizontal motion of its fairlead; to do so the jacket must be accelerated, and the line must supply that acceleration. If the fairlead sways with amplitude X at frequency ω, the force the pair must deliver is

$$F_{\mathrm{dyn}} = M_{\mathrm{i}}\,\omega^2 X, \tag{6.1}$$

where $M_{\mathrm{i}} = 7719\,\mathrm{t}$ is the inertial mass of the jacket, $\omega = 2\pi/T$ is the circular frequency of the excitation and X is the fairlead sway amplitude the jacket is made to follow. The expression is the peak inertia force of a body of mass M_{i} driven harmonically at amplitude X: the acceleration amplitude is $\omega^2 X$, and the force follows from Newton's second law.

A line cannot push, so this force must be delivered as a variation in tension about a pretension large enough that the line never goes slack. The pretension across the pair must therefore be at least F_{dyn}, and the capacity at least $2F_{\mathrm{dyn}}$, since the tension peaks at the pretension plus the variation. Meskers and Dijk [19] applies the same expression to a 3000 t reel and obtains 120 t for 1 m of motion at a 10 s period, so requiring more than 120 t of pretension and more than 240 t of capacity against the 140 t of his installed pair, which is why the mode is not adopted there either.

Table 6.1: Force a fixed-line pair must deliver to make the jacket follow a fairlead sway of amplitude X, $F_{\mathrm{dyn}} = M_{\mathrm{i}}\omega^2 X$, in tonnes. The total pretension must be at least this value to prevent slack, and the total capacity at least twice it.

T [s]	$X = 1\,\mathrm{m}$	$X = 2\,\mathrm{m}$	$X = 3\,\mathrm{m}$
5.5	1016	2032	3048
8.0	485	971	1456
10.0	311	621	932
10.9	260	520	780
15.7	126	252	378

Table 6.1 evaluates Equation (6.1) over the design period range. At the design peak period of 10.9 s, making the jacket follow 1 m of fairlead motion requires 260 t from the pair, and therefore at least 260 t of pretension and 520 t of capacity. The damping arrangement examined in this chapter carries 52 t per line, 104 t across the pair, on a per-line capacity of 102 t, 204 t across the pair. A fixed restraint therefore asks for roughly two and a half times the load on both counts, for a single metre of fairlead motion, and the requirement grows as the square of the frequency: at the short-period end of the design range the same 1 m would need over 1000 t of pretension.

The second ground is frequency, and it is the more damaging. A fixed line acts as a horizontal spring of stiffness $K_{\mathrm{tug}} = E_c A_{\mathrm{eff}}/L_0$ [20, 21]. Added stiffness raises the sway frequency and therefore pulls T_{sway} *down* from 15.70 s.

Meskers and Dijk [19] reports both sides of that shift. In his case it is the intended mechanism: the coupled load roll mode sits at 14 s, inside the swell band, and stiffening the lines moves it below the excitation so that the load travels with the vessel. He also reports the cost at short wave periods the fixed tuggers reduce operability, because the shifted peak lands in the wind-sea range. The present system is the inverse case. The sway period is 15.70 s and the design excitation is at $T_z = 6.6$ s, so there is no band below the excitation to move into: stiffening the plant brings the resonance toward the design sea state, and toward the natural roll period of the vessel, rather than away from them. It is the exact inverse of the detuning strategy of Chapter 5, and the magnitudes in Table 6.1 make it unavoidable as a line stiff enough to restrain a payload of this size is stiff enough to shift its natural period substantially.

6.2.2 Constant-tension tugger

Constant-tension control maintains a fixed tension setpoint, $F_s = F_{\mathrm{pre}}$. The tension velocity slope is then identically zero, $\partial F/\partial v = 0$, so the mode delivers no velocity-proportional damping whatsoever [10, 20, 21]. It is excluded not because it performs poorly but because it does not address the mechanism: it neither dissipates resonant energy nor detunes the plant in a favourable range.

More generally, any law that reacts against *displacement* adds effective stiffness and inherits the frequency objection of Section 6.2.1. What is required is a law that reacts against *velocity* alone adding dissipation without moving the natural period.

6.2.3 Damping tugger

Active damping, introduced for suspended offshore payloads by Meskers and Dijk [19], modulates line tension as a function of pay-out velocity, $F_s = F_{\mathrm{pre}} + f(v)$: tension is increased during pay-out ($v > 0$) and reduced during pay-in ($v < 0$), so that energy is extracted from the payload over each swing cycle. It satisfies the requirement identified above and is adopted for the remainder of this chapter.

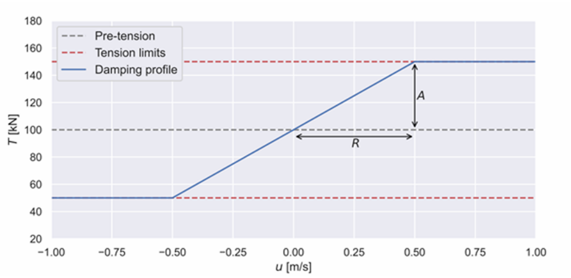

Figure 6.1: Force-velocity relation of the maximum-force setpoint, reproduced from Eerden [10]. The amplitude A and the range R correspond to $\frac{1}{2}(T_{\max} - T_{\min})$ and to the velocity span over which the law is linear.

The selection follows from the mechanism rather than from a preference among three admissible options. A fixed line adds stiffness and moves the sway period in the wrong direction; a constant-tension law adds no damping at all. Only a law reacting against velocity adds dissipation without disturbing the period that Chapter 5 sets out to move. How much that dissipation is worth in the present design envelope, where the excitation sits well below the sway period, is the subject of the remainder of this chapter.

6.3 Control law and actuator model

6.3.1 Setpoint law

The implemented law is the *maximum force setpoint*: a linear force-velocity relation bounded by a maximum and a minimum, evaluated from a measured velocity and applied directly to the load. For the opposing pair, with gains expressed as totals for the system and n lines,

$$T_{s,p}(u) = \mathrm{sat}_{[T_{\min}, T_{\max}]}\left(T_0 \pm \frac{c_{\text{tot}}}{n}\, u\right), \qquad F_{\text{net}} = T_s - T_p = c_{\text{tot}}\, u, \tag{6.2}$$

where T_0 is the pretension and c_{tot} the slope of the force-velocity relation of the system. The pretension is placed midway between the bounds, $T_0 = \frac{1}{2}(T_{\max} + T_{\min})$, matching the linear damper of Sanders [11]; his shifted variants, which place the target at the minimum or the maximum, are reported by him to perform worse as the forward-shifted case extracts too little energy per cycle and the backward-shifted case injects energy through its high pretension.

The velocity u in Equation (6.2) is the *pay-out velocity of the tugger line*, not the velocity of the payload. For a line running from the vessel fairlead at $\mathbf{r}_f$ to the jacket attachment at $\mathbf{r}_j$, of instantaneous length $L = \|\mathbf{r}_f - \mathbf{r}_j\|$ and unit direction $\hat{\mathbf{b}} = (\mathbf{r}_f - \mathbf{r}_j)/L$, it follows that

$$u = \dot{L} = \hat{\mathbf{b}} \cdot (\dot{\mathbf{r}}_f - \dot{\mathbf{r}}_j), \tag{6.3}$$

so u is the component of the fairlead-jacket relative velocity resolved along the live line direction, evaluated with the instantaneous geometry at each time step. Positive u denotes pay-out. This quantity is obtained as the analytic time derivative of the line length and identifies it as the physically measurable input: winches are instrumented to measure pay-in and pay-out velocity, and that velocity is directly related to the payload motion. The two payload and payout velocities coincide only when the line is horizontal and the vessel end is stationary; here the fairlead moves with the vessel and the line direction rotates, so the pay-out velocity carries both the payload motion and the vessel motion at the fairlead

No PID action is applied in the model. In the maximum-force-setpoint formulation the controller output *is* the commanded force: it is computed algebraically from the sensed velocity and applied directly, so there is no tracking error for a PID to act on . This is the configuration Eerden [10] arrives at in his simplified model, and he arrives at it because the detailed model that did include a PID loop failed: gains spanning four orders of magnitude, several tension-setpoint pairs and even removal of the winch limits produced no working controller for the damping mode, which is what prompted the move to the simplified formulation. The abstraction is adopted here because the question is how much damping is worth having on a payload of this size, not how to build the winch controller that delivers it; that a detailed winch model of this kind has already been reported to resist tuning is a further reason not to place one inside the present study.

The lines attach at the height of the jacket centre of gravity. Attaching at the centre-of-gravity height removes the lever for roll, and is the arrangement adopted here.

6.3.2 Actuator model and fairlead

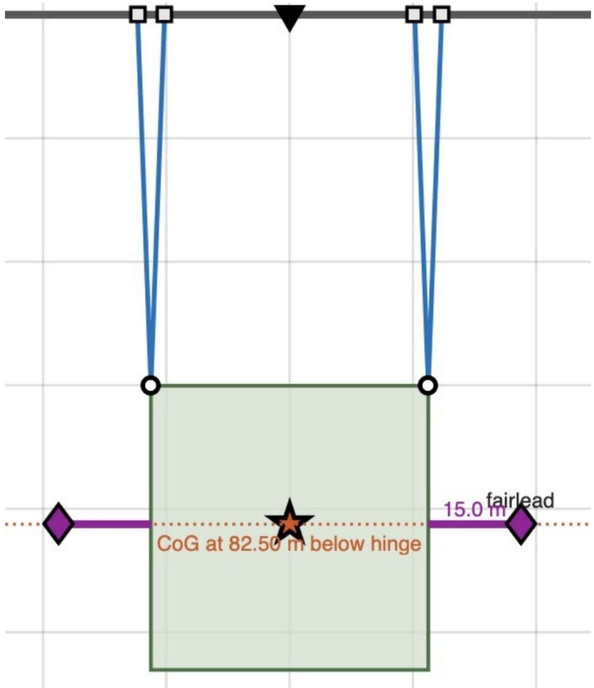

Figure 6.2: Arrangement of tuggers at the pure sway direction of jacket at same level as CoG

The tuggers are modeled as two opposing pairs of forces acting in the direction of relative sway. We assume two fairleads that sit 15 m from the jacket. This configuration is theoretical, as there are no current options to place the tuggers here in reality unless we build large extensions at the PS aft. The winch is modelled as an *ideal force source*: the commanded tension is realised instantaneously and without error, so winch delay $\tau_d = 0$ and the tension is algebraic in the sensed velocity.

Each tugger line runs between two points that both move. The jacket end is attached at CoG level in the pure-sway direction; the ship end is a fairlead fixed to the vessel, with a free line length of 15 m to the jacket. The vessel-side fairlead is not the hoist hinge. It sits 82.50 m below the hinge, and its motion is obtained by transferring the Pioneering Spirit RAOs from the vessel reference point to the fairlead coordinates as a rigid body — the same transfer already applied to the hinge, but with a different lever arm.

Sanders [11] admits exactly two actuator limits: no tension below zero, to avoid slack in the cable, and no tension above the winch capacity. Eerden [10] imposes the same pair. Both are enforced here. The no-slack floor is set at $T_{\min} = 20\,\mathrm{kN}$ rather than zero. The capacity $T_{\max}$ differs between the three configurations examined in Section 6.4 and is given in Table 6.3; the pretension is placed midway between the floor and the capacity in each case.

Two omissions are present in the model. First, no damping-speed limit is imposed. Sanders [11] states explicitly that the maximum pay-out and pay-in speed of the winches is evaluated but *not* set as a hard limit, and a speed threshold is retained here for reporting only. Second, a deadband is not applied. Eerden [10] derives one from an accepted displacement and applies it in his simplified model. Sanders [11], however, tests deadband variants directly and finds their behaviour depends on the loading: they damp a free decay very effectively, but under continuous external forcing the abrupt change in force as the velocity crosses zero degrades motion reduction and generates roll. He also notes that a deadband of this form exists only in the model. The present case is continuous irregular loading, so the deadband is omitted.

The parameters common to all three configurations are collected in Table 6.2. Line axial compliance and winch inertia are not modelled, so the winch, cable and load velocities are equal by construction. The fairlead sits 82.50 m below the hinge that drives the hoist, so the winch and load sensor references differ by vessel roll acting through that lever; with compliance absent, any difference between the two is purely geometric. With a free line length of 15 m the axial stiffness of the tugger line is relatively high.

Table 6.2: Tugger system parameters common to all three configurations. The winch capacity, pretension and damping slope differ between configurations.

Parameter	Symbol	Value
No-slack floor	$T_{\min}$	20 kN
Actuator time constant	τ_d	0 s
Deadband velocity	v_{db}	not applied
Number of lines	n	2
Fairlead offset from centreline	Y_m	37.5 m
Jacket half-width	$W_j/2$	22.5 m

6.3.3 Frequency-domain and time-domain formulations

Two formulations are used. In the frequency domain the tugger is a linear viscous damper of impedance $G(\omega) = \mathrm{i}\omega c$, with no capacity and no slack floor; results obtained there are what the device would give if it never reached a limit. In the time domain the winch limits are active, together with slack-capable hoist wires and the instantaneous line directions. Because the actuator is ideal in both, any difference between them comes from the tension limits and from the nonlinearity of the plant.

Section 6.5 reports the time domain and Section 6.7 the frequency domain. Compared at a fixed wave height the two are close inside the design window, because the tension there rarely reaches a limit: at $T_z = 6.6$ s the reference configuration does not reach one at all. The agreement weakens as the period lengthens at that same wave height at $T_z = 8.5$ s the same configuration sits at its capacity for 16 % of the record. The operability curves of Section 6.7 are a separate case, because the wave height there is not held fixed; the argument for reading them across their whole range is given in that section.

6.4 Three tugger configurations

The control law of Section 6.3 leaves one parameter free: the slope c_{tot} of the force–velocity relation, together with the winch capacity that bounds it. Rather than propose a single tuned device, this chapter takes three configurations and reports what each achieves. The purpose of the study is to provide a way of assessing tugger damping for a given payload and operating envelope, not to specify a winch for one jacket.

The slopes are anchored to the previous systems studied as the PS has no tuggers installed. Both Sanders [11] and Eerden [10] work at approximately 1000 kN·s/m, which is taken as the reference configuration. The low configuration is below that value, and the high one is twice it on winches of twice the capacity. Doubling the slope and the capacity together leaves the pay-out speed at which the law saturates almost unchanged, so the reference and high configurations differ in force authority rather than in where saturation begins.

Table 6.3: The three tugger configurations. Capacity is expressed as a percentage of the payload weight so that it can be compared with other installations.

Configuration	Slope [kNs/m]	$T_{\max}$ [kN]	Pretension [kN]	Saturates at [m/s]	Capacity [% weight]
Low	300	1000	510	3.27	2.7
Reference	1000	1000	510	0.98	2.7
High	2000	2000	1010	0.99	5.3

The capacities are not unusual for this class of operation. Meskers and Dijk [19] reports two 70 t winches restraining a 3000 t load, which is 4.7 % on the same measure. The high configuration, at 5.3 %, is therefore a normally scaled winch rather than an exceptional one.

In every configuration the pretension sits midway between the no-slack floor and the capacity, following Sanders [11]. The lines attach at the height of the jacket centre of gravity, so the tugger force produces no roll moment. No speed limit, deadband or power limit is applied, for the reasons given in Section 6.3.2.

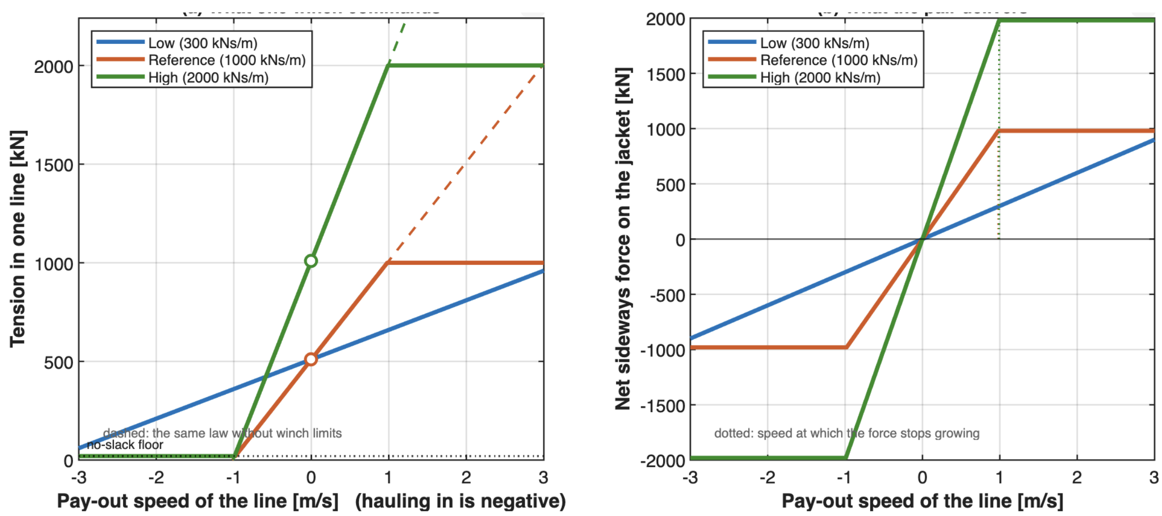

Figure 6.3: Force-velocity relation of the three configurations. Left: the tension commanded in one line against the pay-out speed of that line, linear until the tension reaches a winch limit and flat thereafter; the dashed lines show the same law without limits. Right: the net sideways force the opposing pair applies to the jacket

The net force stops growing at 3.27, 0.98 and 0.99 m/s of pay-out speed for the low, reference and high configurations, and is capped at 980, 980 and 1980 kN respectively. How often the law reaches those limits in realistic seas is reported in Section 6.5.

6.5 Response over an extended sea-state set

At the upper period limit of the beam-sea envelope, H_s = $2.4\,\mathrm{m}$ and T_z = $6.6\,\mathrm{s}$, the jacket already stays inside the 3.0 m clearance budget without any tugger: the largest sway expected in three hours is 2.93 m. A tugger cannot be judged on that cell alone, because the constraint it is meant to relieve is not yet binding. The set of sea states below therefore extends the envelope in both wave height and wave period, and asks what each configuration achieves as the sea gets harder.

Every case uses the nonlinear time-domain form of the model with the winch limits active. The uncontrolled run and the three controlled runs at a given sea state share the same wave record, so the reductions are directly comparable.

Table 6.4: Reduction in the standard deviation of the relative sway achieved by each configuration, ordered by wave period.

H_s $[\mathrm{m}]$	T_z $[\mathrm{s}]$	σ, no tugger [m]	low	reference	high
2.4	6.6	0.785	1 %	10 %	22 %
3.0	6.6	0.983	2 %	10 %	22 %
2.4	7.0	1.141	3 %	15 %	30 %
2.4	7.5	2.762	37 %	48 %	60 %
3.0	7.5	3.637	40 %	49 %	61 %
3.6	7.5	4.252	37 %	46 %	58 %
2.4	8.0	4.270	40 %	53 %	66 %
2.4	8.5	6.196	45 %	56 %	69 %

The reduction rises steadily as the wave period lengthens: from 1 % to 45 % for the low configuration and from 22 % to 69 % for the high one between T_z = 6.6 and 8.5 s. It is set by the period and is almost independent of the wave height: at T_z = $6.6\,\mathrm{s}$ the same percentages are obtained at H_s = 2.4 and 3.0 m. This is the direct consequence of the off-resonance position of the design box established in Section 6.1. At T_z = $6.6\,\mathrm{s}$ only a small share of the spectrum lies near the sway period of 15.7 s, so the oscillatory component the damper acts on is a small part of the total response. As the wave period lengthens toward the sway period the response becomes increasingly dynamic, and the damper has correspondingly more to remove.

Reducing the standard deviation is not the same as making an operation possible. Table 6.5 therefore reports the largest sway expected in three hours against the clearance budget, which is the question that decides whether the transit can proceed.

Table 6.5: Largest sway expected in three hours against the 3.0 m clearance budget. The last column is the fraction of the record for which the high configuration sits at a tension limit

H_s [m]	T_z [s]	none [m]	low [m]	ref. [m]	high [m]	at a limit (high) [%]
2.4	6.6	**2.93**	**2.89**	**2.65**	**2.29**	0
3.0	6.6	3.67	3.61	3.32	**2.86**	0
2.4	7.0	4.23	4.09	3.59	**2.98**	0
2.4	7.5	10.10	6.42	5.27	4.08	1
3.0	7.5	13.29	8.11	6.84	5.25	3
3.6	7.5	15.55	9.86	8.55	6.61	6
2.4	8.0	15.57	9.39	7.47	5.44	3
2.4	8.5	22.56	12.57	10.11	7.02	7

Seven of the eight sea states exceed the budget without a tugger. The low and reference configurations recover none of them; the high configuration recovers two, at $H_s = 2.4\,\mathrm{m}$, $T_z = 7.0\,\mathrm{s}$ and at $H_s = 3.0\,\mathrm{m}$, $T_z = 6.6\,\mathrm{s}$. Both are recovered by a small margin 2 and 14 cm inside the budget and both results come from a single wave realisation.

Where the response is very large the tugger removes a substantial fraction of it without bringing the jacket anywhere near the budget. At $H_s = 2.4\,\mathrm{m}$, $T_z = 8.5\,\mathrm{s}$ the high configuration removes 69 % of the standard deviation, yet the largest expected sway is still 7.0 m. A large percentage reduction on a response that is far outside the budget has no operational meaning, and the two tables must be read together for this reason.

One limit on how far the lower rows can be read should be stated. The three-hour maxima there correspond to swing angles well beyond the small-angle range in which the model was linearised; for those rows the ratios between columns remain informative but the absolute values are indicative only.

Figure 6.4 shows what the tabulated statistics look like as records. It presents a 500 s window from two of the sea states above, taken from the middle of each simulation once the $2300\,\mathrm{s}$ start-up transient has been discarded. These are bounded runs: the winch limits of Section 6.3.2 are active throughout, the hoist wires are free to go slack, and the tugger force is resolved along the instantaneous line directions rather than a fixed one. The actuator itself is ideal, so everything that separates these records from the linear frequency-domain results of Section 6.7 comes from the tension bounds and from the nonlinearity of the plant.

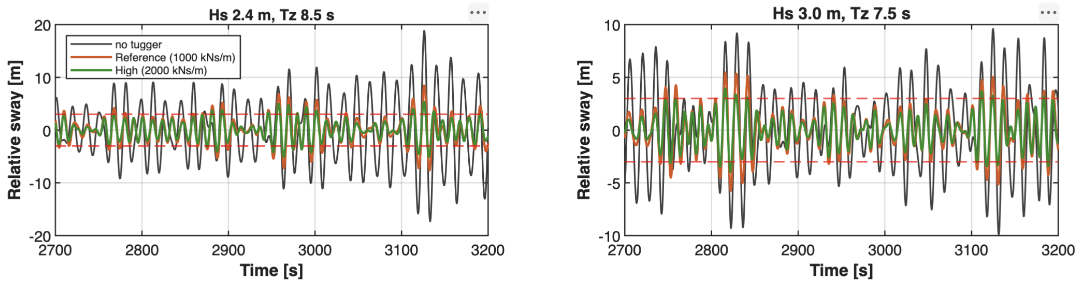

Figure 6.4: Relative sway over a 500 s window for two sea states from Table 6.5, without a tugger and with the reference and high configurations. Dashed lines mark the 3.0 m clearance budget.

Three features are visible. The first is the character of the uncontrolled response. It does not oscillate at a steady amplitude but arrives in bursts: long stretches of moderate motion interrupted by groups of large swings, at $T_z = 8.5\,\mathrm{s}$ reaching almost 20 m in this window alone. This is resonant build-up. Energy accumulates in the lightly damped sway mode over many cycles whenever the wave train happens to contain a run of components near the natural period, and decays slowly afterwards because the structural damping is only 0.98 %. The three-hour maximum is set by these bursts rather than by the typical amplitude, which is why the standard deviation and the most probable maximum tell different stories in Tables 6.4 and 6.5.

The second is what the tugger does to that build-up. The controlled traces are not simply scaled-down copies of the uncontrolled one: the bursts largely disappear, and what remains follows the wave excitation more directly. This is the mechanism of Section 6.1 seen in the time domain. The damper acts on the oscillatory, resonant part of the response and removes most of it, leaving behind the part that mirrors the motion of the suspension point, which added damping barely alters.

The third is that removing the resonant component is not enough. Neither controlled trace stays inside the 3.0 m clearance budget, and the difference between the reference and the high configuration is small compared with the difference either makes against the uncontrolled case. Doubling the damping slope and the winch capacity moves the residual response by much less than the first application of damping did, because what is left after the first application is no longer resonant.

A final point on the extremes. The largest excursion visible in either window is smaller than the three-hour maximum quoted in Table 6.5, because a 500 s sample is unlikely to contain the worst event of a three-hour exposure. The most probable maximum is a statistical extrapolation from the variance and the mean up-crossing period of the record, not the largest value observed in it.

6.5.1 Sensitivity to damping slope

The three configurations are points on a continuum. This section sweeps the damping slope with the winch capacity held at 1000 kN throughout, so the low and reference configurations of Section 6.4 lie on these curves; the high configuration does not, since it is a different device with twice the capacity. The sweep is run in the time domain with both bounds active, so what it shows is the behaviour of a real winch rather than of an ideal damper. Only $T_z = 6.6\,\mathrm{s}$

is a design cell; the two longer periods hold the same wave height and vary the period alone, as a controlled frequency sensitivity.

Table 6.6: Largest sway expected in three hours and the fraction of the record spent at the winch capacity, against the damping slope of the pair, at $H_s = 2.4\,\mathrm{m}$ and a capacity of 102 t.

	$T_z = 6.6\,\mathrm{s}$		$T_z = 8.0\,\mathrm{s}$		$T_z = 8.5\,\mathrm{s}$	
Slope $[\mathrm{kN \cdot s/m}]$	MPM [m]	at cap. [%]	MPM [m]	at cap. [%]	MPM [m]	at cap. [%]
150	2.92	0.0	10.72	0.0	15.08	0.0
300 (low)	2.89	0.0	9.39	0.0	12.57	0.7
600	2.80	0.0	8.14	2.5	10.80	7.2
1000 (ref.)	2.65	0.0	7.47	9.3	10.11	15.7
1500	2.47	0.7	7.14	15.7	9.81	22.1
2000	2.33	2.5	6.99	20.2	9.68	26.6
3000	2.15	6.4	6.85	25.9	9.56	31.5
4500	2.01	11.3	6.76	30.5	9.48	35.3
6178 (crit.)	1.94	14.8	6.71	33.2	9.43	37.6

The two ends of the period range behave differently, and the difference is the central point of this section.

At the design period the benefit does not run out. Increasing the slope from 150 to 6178 kN·s/m, the latter being critical damping for the sway mode reduces the three-hour maximum steadily from 2.92 to 1.94 m, a reduction in the standard deviation rising from 0.5 to 34 % with no sign of a plateau. The winch reaches its capacity for under 1 % of the record up to 1000 kN·s/m and for 15 % even at critical damping.

At the long periods the benefit stops early. At $T_z = 8.5\,\mathrm{s}$ the reduction reaches 56 % at the reference slope and then flattens: increasing the slope sixfold to critical damping buys a further 3 %, moving the three-hour maximum from 10.11 to 9.43 m. The behaviour at 8.0 s is the same.

The cause is visible in the occupancy columns. At the longer periods the payload moves faster, so the pay-out velocity crosses the saturation speed of 0.98 m/s for a substantial part of the record: at $T_z = 8.5\,\mathrm{s}$ the reference configuration already sits at its capacity for 16 % of the time, rising to 38 % at critical damping. Once the tension is on its limit, a steeper force–velocity relation changes nothing . It only moves the saturation point to a velocity the system rarely reaches. The additional slope is therefore spent extending a flat part of the law rather than adding force, and the response stops responding to it.

This is the practical limit on the mitigation, and it is a hardware limit rather than a dynamic one. Recovering the long-period sea states would require a larger winch rather than a steeper law, and Table 6.5 shows what that buys: the high configuration, at twice the capacity, brings the three-hour maximum at $T_z = 8.5\,\mathrm{s}$ from 10.11 to 7.02 m which is still more than twice the clearance budget. Whether a winch large enough to close that gap is worth fitting is a question of cost and deck space rather than of dynamics, and lies outside the scope of this study.

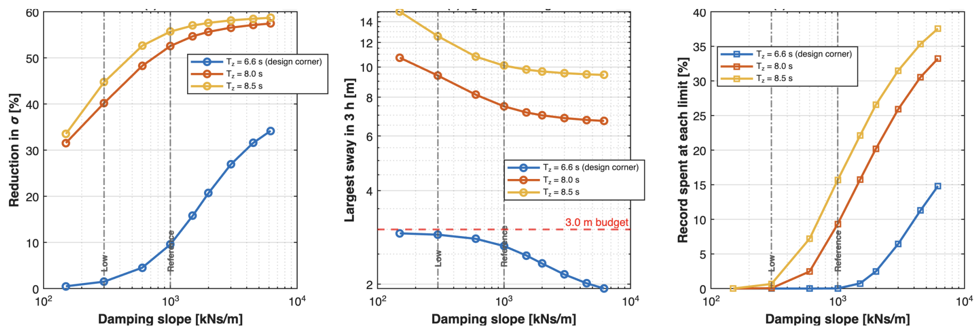

Figure 6.5: Damping slope swept in the time domain at $H_s = 2.4\,\mathrm{m}$ with the winch capacity fixed at 102 t. (a) reduction in the standard deviation of the relative sway; (b) three-hour maximum against the 3.0 m clearance budget; (c) fraction of the record spent at a tension limit.

6.6 Decay behaviour

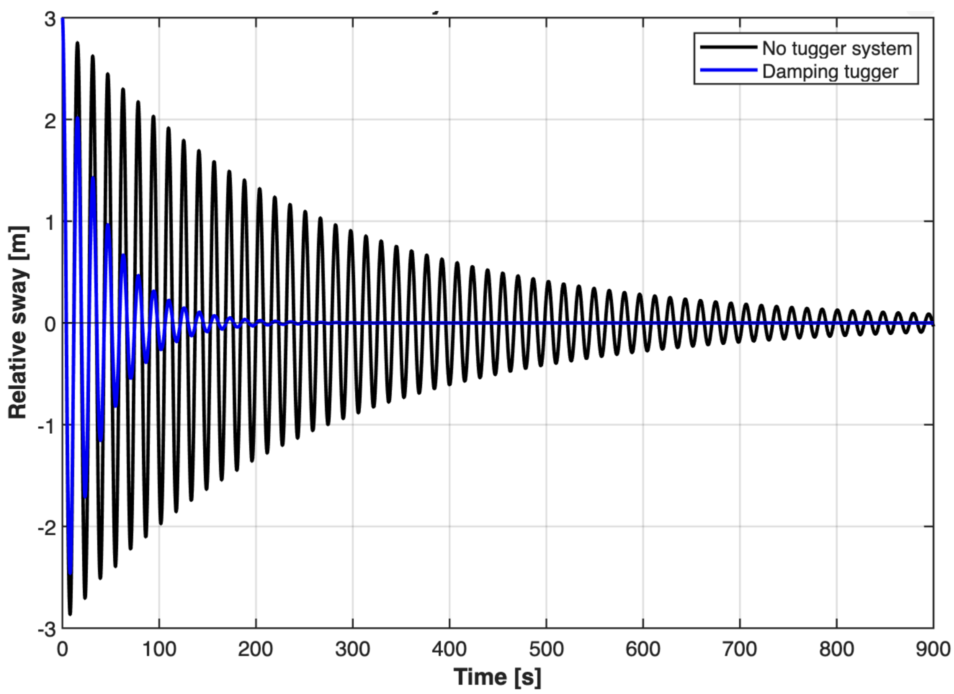

Figure 6.6: Effect on decay test for the 300 kNs/m damping tugger

Releasing the jacket from a 3.0 m offset gives the clearest single demonstration of the added dissipation. The time to decay to 10 % of the release amplitude falls from 581 s without tuggers to 95 s with the low configuaration, a factor of six.

The headline metric is deliberately a time rather than an equivalent damping ratio. Once the tension sits on its rails the dissipation is Coulomb-like and amplitude-dependent, so a fitted ζ_{eq} drifts through the record; and as the mode approaches critical damping the logarithmic decrement ceases to be defined at all. The equivalent damping at the first cycle rises from approximately 1 % to 5.5 %. With the high configuration the decay time falls further, to 17 s.

This test is the resonant limit of the device, and it should be read against Section 6.5 rather than in place of it. A free decay is pure oscillation at the sway period, which is where added damping has most effect, so the same low configuration that shortens the decay by a factor of six removes only 1 % of the standard deviation at the design corner. The two numbers are not in conflict: the decay shows what the damper does to motion at 15.70 s, and the design cell shows how small a share of the response is of that character.

6.7 Operability curves

The time-domain results above give the reduction; this section gives the largest wave height at which the operation can proceed, which is what determines whether a transit goes ahead.

The operability curves are computed in the frequency domain, where the tugger is a linear damper without limits. This is a necessity rather than a preference: the curve is obtained by solving for the wave height at which the three-hour maximum reaches 3.0 m, and that inversion is available in closed form only while the response scales with the wave height. Once the tension bounds are active it does not, and the same curve would require an iterative time-domain search at every period.

The linear form is nevertheless a fair representation here, for a reason that follows from how the curve is built. Every point on it carries the same motion by construction a three-hour maximum of 3.0 m so the response amplitude, and with it the pay-out velocity, is bounded at the same value along the whole curve. The saturation reported in Section 6.5.1 occurs at a fixed wave height of 2.4 m while the period is lengthened, where the uncontrolled response reaches 10 to 22 m, three to seven times the budget; those velocities are not reached anywhere on the operability curve. At the current design upper limit this is verified directly: at $T_z = 6.6\,\mathrm{s}$ the reference configuration does not reach a tension limit for any part of the record (Table 6.6).

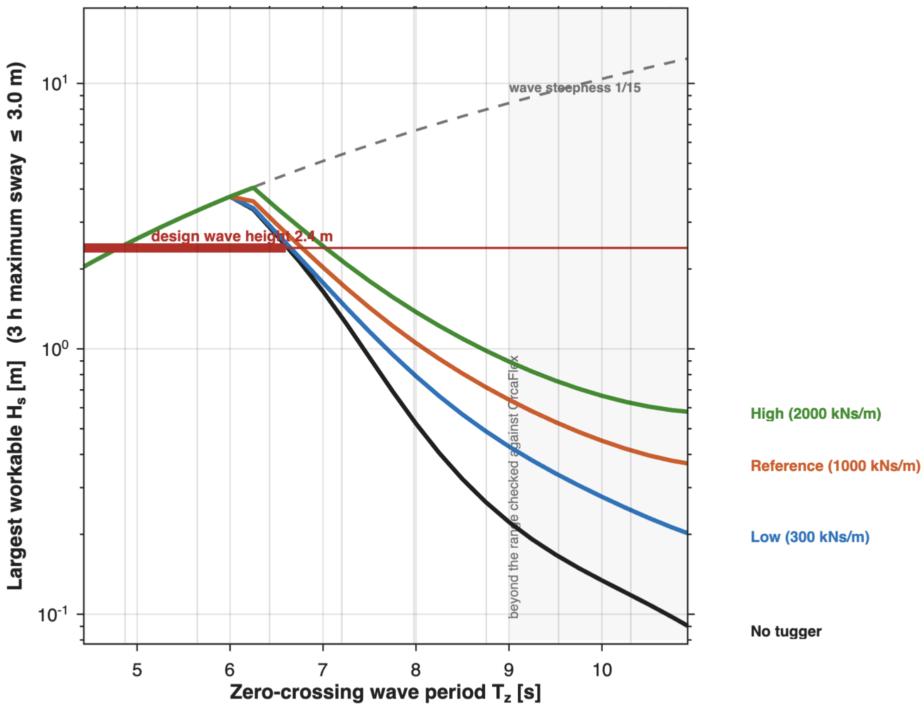

Figure 6.7: Largest workable significant wave height against zero-crossing wave period for the three configurations. The dashed grey line is the wave steepness ceiling. The horizontal red line is the current wave height, with the present operating window marked in bold.

Table 6.7: Workable wave-period window at the design wave height of 2.4 m. The lower end is set by the wave steepness ceiling and is the same for every configuration.

Configuration	to T_z [s]	gained [s]
No tugger	6.63	—
Low (300 kN·s/m)	6.66	+0.03
Reference (1000 kN·s/m)	6.78	+0.15
High (2000 kN·s/m)	7.03	+0.40

Two features of this table matter more than the gains themselves.

First, the uncontrolled configuration is workable to $T_z = 6.63$ s, and the operator's present envelope at this heading ends at $T_z = 6.6$ s. The two agree to within the resolution of the calculation. The 3.0 m budget was set by Allseas based on the operational guidelines.

Second, the lower end of the window is identical for all four configurations. At periods below about 5 s the wave steepness ceiling of $1/15$ limits the sea state before jacket motion does. No tugger, of any size, can move a limit that is not about motion.

The gains at the upper end are modest: 0.03, 0.15 and 0.40 s for the three configurations. Two things make them small, and only the first belongs to the tugger. The reference configuration removes about 10 % of the sway at the design corner (Table 6.4), and that translates directly into workable height: at $T_z = 6.50$ s it raises the limit from 2.68 to 2.95 m, also 10 % (Table 6.9). The second is the shape of the curve: near the design corner the workable wave height falls by roughly 2.2 m for every second of period, so 10 % of 2.68 m converts to only about 0.12 s. The rigging change raises the same limit to 3.77 m, 41 %, and returns 0.50 s through the same conversion. A steep operability curve, not the damper alone, is what compresses both gains into tenths of a second.

6.8 Hoist optimisation combined with damping tugger

Chapter 5 changed the rigging geometry and this chapter has added damping. Both act on the same motion, so this section places them on the same axes, together with the case where both are applied, and establishes which of the two should be preferred.

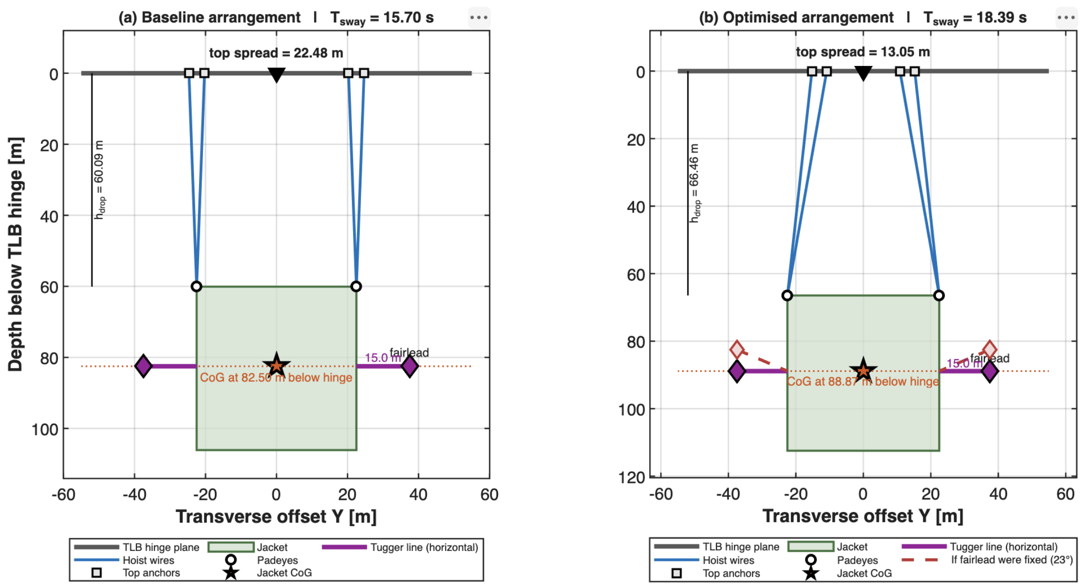

Figure 6.8: The two rigging combined with tugger arrangements compared in this section.

All four cases use the same plant, the same 3.0 m clearance budget and the same frequency-domain calculation, so the comparison is like for like. The tugger is the reference configuration of 1000 kN·s/m. The optimised geometry is the top spread and hoist drop identified in Chapter 5, shown in Figure 6.8.

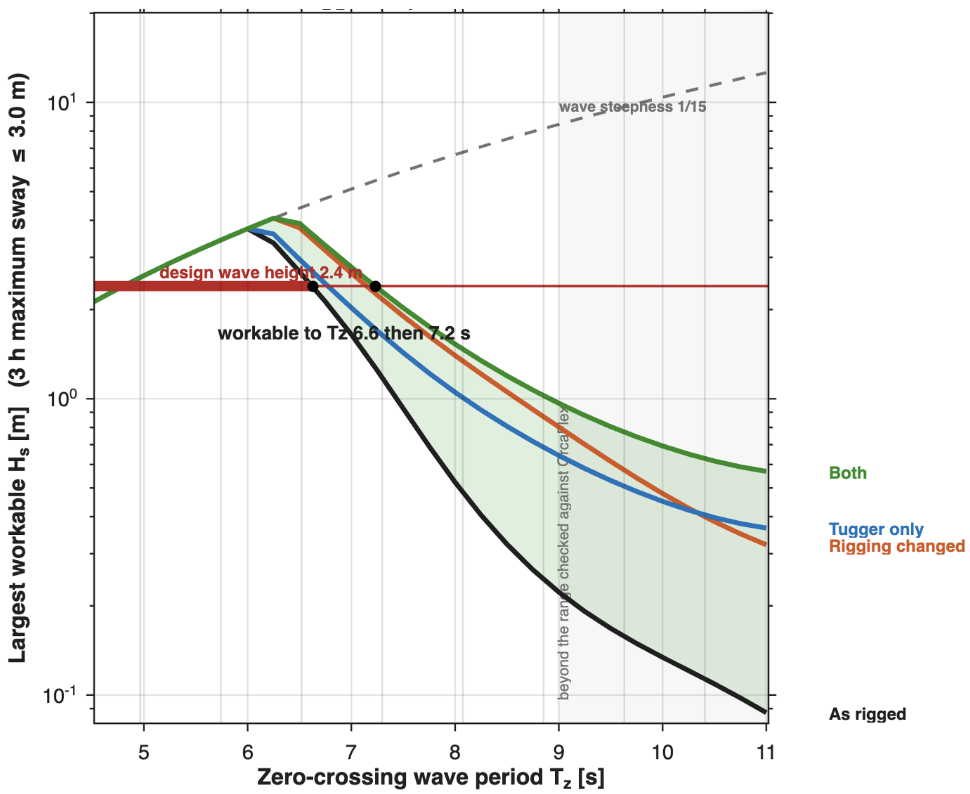

Figure 6.9: Largest workable wave height for the four cases. The shaded area between the as-rigged and combined curves is the total gain available from both mitigations together. The dashed grey line is the wave steepness ceiling.

Table 6.8: Workable wave-period window at the design wave height of 2.4 m. The lower end is set by the wave steepness ceiling and is common to all four cases.

Case	to T_z [s]	gained [s]
As rigged	6.63	—
Tugger only	6.78	+0.15
Rigging changed	**7.15**	**+0.53**
Both	7.23	+0.60

The rigging change is worth about three and a half times the tugger: 0.53 s of additional workable period against 0.15 s. Applying both gives 0.60 s, which is less than the 0.68 s that adding the two separate gains would suggest.

That shortfall is not an inconsistency but a consequence of how the two act. Lengthening the sway period moves the excitation further above the sway resonance, so less of the spectrum overlaps it and a smaller share of the response is damping-controlled. A damper acts on velocity, so it has less to work with once that has happened. The two mitigations address the same motion, and their effects therefore overlap rather than add. The same mechanism that makes the tugger weak inside the present design envelope, set out in Section 6.1, makes it weaker still once the rigging has been detuned.

Table 6.9: Largest workable wave height across the period range. Where a value equals the steepness ceiling, the sea rather than the jacket sets the limit and no mitigation can change it.

T_z [s]	as rigged [m]	rigging [m]	tugger [m]	both [m]	steepness ceiling [m]
6.00	3.75	3.75	3.75	3.75	3.75
6.25	3.35	4.07	3.60	4.07	4.07
6.50	2.68	3.77	2.95	3.91	4.40
6.75	2.13	3.16	2.44	3.29	4.74
7.00	1.65	2.65	2.03	2.78	5.10
7.25	1.25	2.24	1.69	2.36	5.47
7.50	0.93	1.90	1.43	2.03	5.85
8.00	0.52	1.40	1.05	1.53	6.66

Two regimes are visible in this table.

At $T_z = 6.00\,\mathrm{s}$ all four cases give the same limiting wave height, and that height equals the steepness ceiling. Over most of the present design window the sea itself is the binding constraint: waves steeper than approximately $1/15$ do not occur, so the limit is reached before jacket motion becomes critical. Neither mitigation can improve operability there, and no larger winch would change that.

Above about 6.6 s the limit becomes motion-governed and the four cases separate. At $T_z = 7.00\,\mathrm{s}$ the as-rigged arrangement is limited to 1.65 m, the tugger raises this to 2.03 m, the rigging change to 2.65 m and both together to 2.78 m. The ordering is the same at 8.00 s. In the region where a mitigation can help at all, the rigging change consistently delivers more than the tugger, and the two together deliver only a little more than the rigging change alone.

This is the central result of the comparison, and it is worth stating plainly. The rigging change requires no additional equipment, no control system and no crew intervention; it is a matter of how the jacket is rigged before the transit begins. The tugger requires two winches of 100 t capacity, their control system and their maintenance, and delivers less. For the present envelope the rigging change is therefore the mitigation to adopt.

The tugger is not thereby dismissed. Its effectiveness rises steeply with wave period, as Section 6.5 showed, so a payload with a shorter sway period, or an operation extending into longer waves than the present design table allows, would place the excitation closer to resonance and give the damper considerably more to remove. What the present results establish is the condition under which a damping tugger is worth fitting, rather than a conclusion that it never is.

6.9 Validity and limitations

This is a numerical investigation. No tugger system exists on the current JLS arrangement, so there is nothing to validate the control result against; the study is therefore a feasibility assessment rather than a verified design.

The principal limitations, each a scoping decision rather than an oversight, are as follows.

- **Ideal force source.** No actuator lag, no sensor error, tension algebraic in the sensed velocity, following Sanders [11], who assumes a perfectly functioning winch with no error and no delay. The framework accepts a first-order lag by setting $\tau_d > 0$.
- **No line compliance or winch inertia.** Winch, cable and load velocities are equal by construction, following Eerden [10], who adopts the same assumption. Relaxing it would introduce a sensed velocity distinct from the load velocity.
- **Relative sway geometry.** The geometric arrangement assumed to calculate the relative sway including the fairlead offsets, the idealized line angles, and the attachment at the centre-of-gravity height to eliminate roll coupling is adopted strictly for the purpose of this numerical investigation. A physical implementation would require adapting the assumed geometric relations and payout velocity calculations to the actual as-built rigging and operational deck constraints.

6.10 Conclusions

1. A velocity-proportional law is the only admissible tugger mode for this system. Fixed lines require roughly two and a half times the pretension of the damping arrangement and add stiffness that shortens the sway period toward the wave band, which is the opposite of what Chapter 5 sets out to achieve. Constant tension provides no damping by construction.
2. The current design envelope sits well below the sway period of the jacket. Only part of the response there is oscillation at the sway period — the part added damping can reduce — and that part grows as the wave period lengthens. At the current design transportation state T_z limit for beam seas, the reduction in the standard deviation of the sway is 1 % for the low configuration, 10 % for the reference and 22 % for the high; at $T_z = 8.5\,\mathrm{s}$, approaching the sway period, the same three give 45 %, 56 % and 69 %.
3. How much damping is worth fitting depends on the wave period. At the design period the benefit does not run out: raising the slope from 150 kN·s/m to critical damping reduces the three-hour maximum steadily from 2.92 m to 1.94 m . At longer wave periods, winch saturation dictates the limit of active damping performance. At $T_z = 8.5\,\mathrm{s}$, the reduction reaches 56 % and then flattens aggressively, gaining a further 3 % for a sixfold increase in damping. This occurs because the winch already spends 16 % of the record at its capacity, rising to 37.6 % at critical damping. Achieving further sway reduction in these boundary states requires a substantially larger winch rather than a steeper law, making it an economic and deck space question rather than a dynamic one.
4. Rigging geometry and tugger damping are sub-additive, and the rigging change dominates. At the design wave height the optimised geometry extends the workable period by 0.53 s against 0.15 s for the reference tugger, and applying both gives 0.60 s rather than the 0.68 s that independence would suggest. Lengthening the sway period moves the excitation further from the sway resonance, leaving a smaller damping-controlled share of the response, which is precisely the condition in which a damper has least to remove.
5. For the present envelope the recommendation is therefore to mitigate through the rigging geometry, which requires no additional equipment. The margin is wider than the

curves alone suggest, because the comparison is generous to the tugger at every step. The arrangement assumed here places the fairleads 15 m from the jacket at centre-of-gravity height, which is the most favourable geometry available the line acts entirely in the sway direction and exerts no roll moment and which the present aft arrangement of the vessel cannot provide without structural extension. The winch is an ideal force source with no lag and no sensor error, and the operability curves impose no tension limit at all. Even on those terms the damper extends the workable period by 0.15 s against 0.53 s for a change in how the jacket is rigged. A realisable fairlead would sit further outboard and lead at an angle, reducing the horizontal component of the line force and reintroducing a roll lever, so the figures above are an upper bound on what a fitted system would deliver and the conclusion holds with room to spare.

6. The tugger is not thereby dismissed. Its effectiveness rises steeply with wave period, so a payload with a shorter sway period, or an operation extending into longer waves than the present design table allows, would place the excitation closer to resonance and give the damper considerably more to remove. What these results establish is the condition under which a damping tugger is worth fitting, not a conclusion that it never is.

Chapter 7

Conclusions and Recommendations

7.1 Conclusions

This work set out to explain why a suspended jacket sways during transit on the Pioneering Spirit and to establish what can be done about it. The mechanism is now clear. For a typical jacket we studied, the rigged payload acts as a lightly damped pendulum with a sway period of $15.70\,\mathrm{s}$ and a structural damping of approximately $1\,\%$. Because the current beam-sea design states excite the system below this period, the response consists of two parts. Most of the motion follows the suspension point at the wave period, while a smaller component oscillates at the sway period. This second component arrives in bursts and dictates the three-hour maximum against which clearance is judged. Yaw is not independently excited but is carried along by sway through the off-diagonal stiffness of the asymmetric four-wire rigging.

Sub-question 1 asked how a reduced-order model could be constructed that captures the dominant swell-induced dynamics with sufficient fidelity for control design. A four-degree-of-freedom model is sufficient. Vessel motion enters as a prescribed kinematic boundary input, transferred from the vessel centre of gravity to the tilting lift beam hinge, and the restoring and base-excitation matrices follow directly from the as-rigged geometry. The resulting model reproduces the free-decay behaviour, transfer functions, and irregular-sea response of the nonlinear OrcaFlex reference, and predicts the three-hour maximum of a measured transit to within 3 %. Outside the design envelope it consistently over-predicts in one direction, making it a reliable and conservative basis both for rapid design iteration and for the control synthesis.

Sub-question 2 asked how the hoist-arrangement parameters influence the dominant natural periods and how far they can be tuned passively. The sway response is governed by geometric restoring stiffness, which is highly sensitive to the top spread and, as in any pendulum, to the drop height. Narrowing the top spread from $22.48\,\mathrm{m}$ to $13.05\,\mathrm{m}$ and paying out the available drop lengthens the sway period from 15.70 s to 18.39 s. At the most severe verification state this passive detuning halves the relative sway and carries the relative yaw down with it; in operability terms it extends the workable period at the design wave height by 0.53 s. The optimum sits on the boundary of the geometric constraint set, and further detuning would continue to reduce the response. The limit is therefore structural, not dynamic: how far the rigging can be moved depends on the loads it imposes on the jacket.

Sub-question 3 asked whether an active tugger-line controller can increase workability beyond what the hoist adjustment already achieves. Within the present envelope it cannot, by a wide margin. A velocity-proportional law is the only admissible control mode, as fixed lines introduce unwanted stiffness and constant tension provides no damping by construction. Because the jacket is driven well above its own sway natural frequency, most of the relative motion is inertia-controlled: the jacket stays nearly stationary in space while the suspension point moves beneath it, and that part of the response is largely insensitive to added damping. Only the component near the sway period is damping-controlled, and at the design corner it accounts for roughly a fifth of the variance. Lengthening the wave period brings more of the spectrum's low-frequency tail onto the resonance, which is why the same device removes $10\,\%$ of the sway standard deviation at T_z = $6.6\,\mathrm{s}$ and $56\,\%$ at 8.5 s.

The comparison is also insightful. A change in how the jacket is rigged requires no additional equipment, no control system, and no crew intervention, yet it significantly outperforms a tugger system that requires all three. Taken together, the answer to the main research question is 0.60 s of additional workable period at the design wave height of 2.4 m, extending the limit from 6.63 s to 7.23 s in T_z, of which the rigging change alone delivers 0.53 s. For the present envelope, rigging geometry optimisation is the recommended mitigation strategy. However, active damping is not entirely dismissed. Its effectiveness rises steeply as the excitation approaches the sway period. For a payload with a shorter natural period, or for operations extending into longer waves than the present design table allows, active tuggers remain a viable and necessary dynamic solution.

7.2 Recommendations

For the present operation, the primary recommendation is to mitigate motion through the rigging geometry, subject to one structural condition. The optimisation carried out here is strictly a dynamic study. It constrains lead angles, splash clearance, and padeye positions, but it does not resolve how the changed sling inclination impacts the static tension distribution or the load introduction at the lift points. Lower motions reduce the inertial contribution while the altered geometry changes the static contribution, meaning the net effect on structural utilisation cannot be read from the motion results alone. A dedicated structural assessment of the jacket and its rigging must precede any operational adoption. Beyond that, the framework is readily generalisable. Its inputs (payload mass and inertia, padeye layout, vessel response operators, and sea states) are all known before a transit. Replacing the conservative envelope used in this study with the wave scatter of a planned voyage will return the best rigging for that specific route. The screening cost is low enough for this to become a standard pre-transit engineering activity.

On the active side, the present result establishes when a damping tugger is worth fitting rather than how one should be built. The fairleads assumed here sit 15 m from the jacket at centre-of-gravity height, which is the most favourable geometry available and one the present aft structure cannot provide; repeating the assessment at fairlead positions the vessel can actually offer would show how much of the benefit survives a realisable lead angle. A more detailed actuator, carrying winch lag, line compliance and inertia, would then show whether the damping remains effective once the sensed velocity no longer equals the load velocity. Commanded tension and pay-out velocity are both already available from the time-domain records, so their product gives the power the winch must deliver and absorb. Reporting peak and mean power alongside the motion reduction would allow a candidate winch to be screened against the installed power and deck space on the vessel, turning the choice of tugger into an engineering selection rather than a dynamic question.

The more substantial recommendation concerns the optimisation objective itself. This study minimised relative sway on the grounds that yaw follows it and clearance is the binding criterion. The verification proved this assumption holds, but only for this specific payload and only with the block connection system engaged. Roll, which was never constrained, was the one degree of freedom to move in the unfavourable direction. A sway-only objective offers no warning if roll becomes critical on a different jacket. Carrying the rotational motions as explicit constraints is the minimum required extension.

The most valuable extension, however, is to change the target variable entirely. Motion is ultimately a proxy. What actually limits the operation is clearance on one side and the loads carried by the rigging on the other. Every degree of freedom (sway, yaw, and roll) expresses itself in the line tensions. An objective written on peak dynamic tension and structural utilisation would capture the coupled behaviour without having to enumerate it, unifying the motion and structural assessments that this study kept separate. The prerequisite for this is a validated analytical model of the wire tension. While the reduced-order model already carries the wire stiffness and exact geometry, the tension channel is where a linearised model is most exposed to unmodelled effects like slack and snap events, geometric stiffening at large angles, and the discrete contact mechanics of the block connection system. Establishing that a tension model of this kind can be trusted against ORCAFLEX and full-scale measurement is the natural next step to turn the present screening method into a comprehensive rigging design tool.

References

[1] J. M. J. Journée and W. W. Massie, *Offshore Hydromechanics*. Delft University of Technology, 2001.

[2] K. Hasselmann et al., "Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP)," *Deutsche Hydrographische Zeitschrift, Reihe A*, vol. 12, 1973.

[3] C. Lucas and C. Guedes Soares, "On the modelling of swell spectra," *Ocean Engineering*, vol. 108, pp. 749–759, 2015. DOI: `10.1016/j.oceaneng.2015.08.017`.

[4] T. de Vlieger, "Theory related to reducing payload motion during offshore lifting operations," M.S. thesis, Delft University of Technology, 2022.

[5] DNV, "DNV-RP-H103: Modelling and analysis of marine operations," Det Norske Veritas, Tech. Rep., 2011.

[6] A. Mackojc and B. Chilinski, "Preliminary modelling methodology of a coupled payload–vessel system for offshore lifts of light and heavyweight objects," *Bulletin of the Polish Academy of Sciences: Technical Sciences*, vol. 70, no. 1, e139003, 2022. DOI: `10.24425/bpasts.2021.139003`.

[7] M. Sun et al., "Multi-cable anti-swing system for cranes subject to ship excitation and wind disturbance: Dynamic analysis and application in engineering," *Ocean Engineering*, vol. 281, p. 114 631, 2023. DOI: `10.1016/j.oceaneng.2023.114631`.

[8] J. K. Woodacre, R. J. Bauer, and R. A. Irani, "A review of vertical motion heave compensation systems," *Ocean Engineering*, vol. 104, pp. 140–154, 2015. DOI: `10.1016/j.oceaneng.2015.05.004`.

[9] Y. Lian et al., "Experimental investigation on service safety and reliability of full-scale HMPE fiber slings for offshore lifting operations," *Ocean Engineering*, vol. 287, p. 115 858, 2023.

[10] K. Eerden, "Motion control systems for offshore heavy lifting operations," In collaboration with Heerema Marine Contractors, M.S. thesis, Delft University of Technology, Oct. 2023.

[11] N. Sanders, "Work-ability increase using tugger control," In collaboration with Boskalis, M.S. thesis, Delft University of Technology, Aug. 2019.

[12] B. Chilinski, A. Mackojc, R. Zalewski, and K. Mackojc, "Proposal of the 3-DOF model as an approach to modelling offshore lifting dynamics," *Ocean Engineering*, vol. 203, p. 107 235, 2020. DOI: `10.1016/j.oceaneng.2020.107235`.

[13] D. F. Domingos, P. Wellens, and J.-W. van Wingerden, "Frequency-domain framework for floating installation of wind-turbine towers," *Ocean Engineering*, vol. 297, p. 116 952, 2024. DOI: `10.1016/j.oceaneng.2024.116952`.

[14] K. J. Jensen, M. K. Ebbesen, and M. R. Hansen, "Anti-swing control of a hydraulic loader crane with a hanging load," *Mechatronics*, vol. 77, p. 102 599, 2021. DOI: `10.1016/j.mechatronics.2021.102599`.

[15] Makita, *Principle of a tuned mass damper*, Online, 2023.
[16] P. Atzampou, "Contactless control of offshore wind turbines during installation: A motion control technique based on magnetic interaction," Ph.D. dissertation, Delft University of Technology, 2025. DOI: 10.4233/uuid:8fd5f93a-7ab5-4228-8f58-cf0c4c6f9ff5.
[17] P. Atzampou, P. C. Meijers, A. Tsouvalas, and A. V. Metrikine, "Contactless control of suspended loads for offshore installations: Proof of concept using magnetic interaction," *Journal of Sound and Vibration*, vol. 575, p. 118 246, 2024. DOI: 10.1016/j.jsv.2024.118246.
[18] L. Cao, H. Li, and N. Chen, "Anti-swing control of a floating crane system with underactuated characteristics," *Ocean Engineering*, vol. 217, p. 107 936, 2020.
[19] G. Meskers and R. van Dijk, "A damping tugger system for offshore heavy lifts," in *Proceedings of the ASME 31st International Conference on Ocean, Offshore and Arctic Engineering (OMAE2012), Volume 1: Offshore Technology*, American Society of Mechanical Engineers, 2012, pp. 315–323. DOI: 10.1115/OMAE2012-83334.
[20] H. Lageveen, "Modelling and steering of tuggerwinch," In collaboration with Heerema Marine Contractors, M.S. thesis, Delft University of Technology, Apr. 2014.
[21] M. T. Bron, "Model-free positioning control of a payload by use of wires," In collaboration with Boskalis, M.S. thesis, Delft University of Technology, Norwegian University of Science, and Technology, Aug. 2022.
[22] Z. Ren, A. S. Verma, B. Ataei, K. H. Halse, and H. P. Hildre, "Model-free anti-swing control of complex-shaped payload with offshore floating cranes and a large number of lift wires," *Ocean Engineering*, vol. 228, p. 108 868, 2021. DOI: 10.1016/j.oceaneng.2021.108868.
[23] J. Yang, E. M. He, and Y. Q. Hu, "Dynamic modeling and vibration suppression of floating crane with tuned mass damper," *Ocean Engineering*, vol. 243, p. 110 196, 2022.
[24] Z. Jiang, Z. Gao, Z. Ren, Y. Li, and S. Mukerji, "Dynamic response analysis of a wind turbine blade installation vessel under wave and wind excitations," *Marine Structures*, vol. 58, pp. 412–428, 2018. DOI: 10.1016/j.marstruc.2017.05.011.
[25] R. A. McKenzie and R. A. Irani, "Motion compensation for maritime cranes during time-varying operations at the pendulum's natural frequency," *Mechanism and Machine Theory*, vol. 168, p. 104 573, 2022. DOI: 10.1016/j.mechmachtheory.2021.104573.
[26] W. Guachamin Acero, L. Li, Z. Gao, and T. Moan, "Methodology for assessment of the operational limits and operability of marine operations," *Ocean Engineering*, vol. 125, pp. 308–327, 2016. DOI: 10.1016/j.oceaneng.2016.08.015.
[27] M. A. Kist, "Development of sway reduction control for the jacket lift system of pioneering spirit," Delft Center for Systems and Control, M.S. thesis, Delft University of Technology, 2018.
[28] A. K. Chopra, *Dynamics of Structures: Theory and Applications to Earthquake Engineering*, 5th ed. Hoboken, NJ: Pearson, 2017.
[29] DNV GL, "DNVGL-RP-N103: Modelling and analysis of marine operations," DNV GL, Tech. Rep., 2017.

Reflection on the Use of Artificial Intelligence

Approach to AI Use

Throughout this thesis, I used artificial intelligence (AI) as a productivity and technical-support tool rather than as a source of engineering knowledge. AI was not used to formulate theoretical models, interpret numerical simulation results, evaluate structural dynamics, or draw final conclusions. Instead, all engineering arguments, discussions, and conclusions presented in this thesis were developed strictly from my own analytical derivations, numerical modelling, and the original academic literature.

My use of AI was therefore restricted to supportive tasks that improved efficiency without replacing engineering judgement or critical evaluation.

Software, Coding, and Reporting Support

Large language models (LLMs), primarily Gemini Pro, were utilized to assist with the technical execution and documentation of this research. Specifically, AI was used to support the following areas:

- **Modelling and Simulation:** Clarifying software mechanics, resolving technical doubts, and troubleshooting syntax within ORCAFLEX and MATLAB.
- **Data Processing:** Debugging MATLAB scripts used for data handling, post-processing of simulation outputs, and the generation of technical figures.
- **Reporting and Writing:** Refining language, improving sentence flow, and structuring arguments for clarity during the drafting phase.
- **Document Preparation:** Generating and modifying LaTeX code to manage page layouts, format large tables, typeset equations, and resolve compilation errors in Overleaf.

By relying on AI for software debugging and reporting logistics, I was able to dedicate significantly more time and focus to the core structural engineering analysis and motion control research.

Appendices

Appendix A

Mass and Stiffness of the Reduced-Order Model

This appendix supports the modelling decisions of Chapter 3. Section A.1 lists the numerical matrices. Section A.2 sets out how the inertial mass carried by the sway mode is bounded and identified, and why it is not computed directly. Section A.3 shows that the sway period is effectively independent of the hoist-wire stiffness, which is the reason the optimisation of Chapter 5 treats only geometric variables.

A.1 Numerical mass and stiffness matrices

```
============================================================
STEP 3: NUMERICAL MATRICES (4x4)
============================================================
Mass Matrix M (4x4):
        7719           0           0           0
           0        7719           0           0
           0           0     4550000           0
           0           0           0     3286000
Stiffness Matrix K (4x4):
   1.0e+08 *

    0.0000         0    0.0004    0.0005
         0    0.0023         0         0
    0.0004         0    1.2099   -0.0205
    0.0005         0   -0.0205    0.1024
Base Excitation K_base (4x4):
   1.0e+08 *

   -0.0000         0    0.0002   -0.0005
         0   -0.0023         0         0
   -0.0004         0   -1.1794    0.0205
   -0.0005         0    0.0653   -0.1024
Diagonal stiffness terms:
  K(Sway,Sway) = 1.559e+03 kN/m
  K(Heave,Heave) = 2.339e+05 kN/m
  K(Roll,Roll) = 1.210e+08 kN m/rad
  K(Yaw,Yaw) = 1.024e+07 kN m/rad
Cross-coupling stiffness terms (non-zero = coupled modes):
  K(Sway,Roll) = +3.695e+04
  K(Sway,Yaw) = +5.425e+04
  K(Roll,Yaw) = -2.052e+06
Stiffness asymmetry: 0.000e+00  (should be ~0)
Symmetry diagnostic (should hold exactly for translational DOFs):
  K_base(1,1) + K(1,1) = 0.000e+00   (relative-coord sway,  must be 0)
  K_base(2,2) + K(2,2) = 0.000e+00   (relative-coord heave, must be 0)
  K_base(3,3) + K(3,3) = 3.048e+06   (rotational, expected nonzero)
  K_base(4,4) + K(4,4) = 0.000e+00   (rotational, expected nonzero)
```

Figure A.1: Numerical mass matrix $\mathbf{M}$, structural stiffness matrix $\mathbf{K}$ and base-excitation stiffness matrix $\mathbf{K}_b$ of the 4-DOF reduced-order model, evaluated at the baseline hoist arrangement.

A.2 Inertial mass of the sway mode

A.2.1 Why the participating mass is not computed directly

The sway mode does not move the payload alone. The four 190 t MH blocks sit at the junction of the steel and Dyneema segments, partway up the suspension, and the hoist wires themselves have mass. Both move with the mode and therefore contribute to its kinetic energy. Chapter 3 gives the kinematic lumping that converts this contribution into an equivalent mass at the jacket centre of gravity, scaled by the square of the lever-arm ratio $s_{\text{block}}/L_{\text{total}}$.

Two obstacles prevent that expression from being used as the model parameter.

The first is that the lever-arm ratio is a property of the *arrangement*, not of the payload. At the baseline geometry the blocks sit 38.61 m below the top anchor on a suspension of 62.21 m, giving a ratio of 0.62 and hence a participation factor of 0.385. Every candidate geometry in the design sweep of Chapter 5 moves that ratio, and the optimum moves it substantially: the drop increases from 60.09 to 66.46 m while the top spread narrows from 22.48 to 13.05 m. A directly computed added mass would therefore be a different number for each of the 441 geometries examined, which defeats the purpose of a frozen linear plant.

The second is that the expression assumes the blocks follow the pendular velocity profile freely, and they do not. The Block Connection System—the constant-tension winches and nonlinear damper links described in Section 3.1—ties the blocks to the lift beams. The free-decay evidence of Appendix B.4 shows this restraint is dynamically significant: removing the BCS lengthens the decay period from 15.41 to 15.71 s, so the system carries measurable additional stiffness through that path. A restraint that alters the stiffness also alters how much of the block motion is transmitted, and the fraction is not known a priori.

A.2.2 Bounding the participating mass

The two limiting assumptions bracket the answer. If the blocks are taken to participate fully, the kinematic lumping of Chapter 3 adds 290 t. If they are taken to be fully restrained by the BCS, the added mass is zero and the payload mass stands alone. Evaluating the sway period under each assumption, against the measured free-decay period of the bare arrangement, gives Table A.1.

Table A.1: Sway period under the two limiting assumptions about MH block participation, and under the identified value. The reference is the measured free-decay period of the bare arrangement, 15.708 s, since the reduced-order model contains no BCS.

Assumption	Basis	M_{kin} [t]	T_{sway} [s]	Error
Blocks fully participating	kinematic lumping	7940	15.92	1.4 %
Identified	**bare free decay**	**7719**	**15.70**	**0.00 %**
Blocks fully restrained	payload mass only	7650	15.63	–0.44 %

The identified value lies between the bounds, and much closer to the restrained limit: the 69 t identified is 24 % of the 290 t that free participation would give. This is physically consistent with the role of the BCS. The blocks are not free bodies hanging on the suspension; they are tethered to the lift beams by winches and damper links, and follow the jacket's swing only partially.

A.2.3 Identifiability

Identifying a mass from a period is only legitimate if the period is in fact sensitive to mass and insensitive to the other unknowns. Both hold here. A sweep of the mass and the hoist-wire stiffness about the nominal point gives

$$\frac{\partial T_{\mathrm{sway}}}{\partial M_{\mathrm{kin}}} = 0.099\,\mathrm{s}\text{ per }100\,\mathrm{t}, \qquad \frac{\partial T_{\mathrm{sway}}}{\partial k} = -0.007\,\mathrm{s}\text{ per }10\,\mathrm{MN/m}, \tag{A.1}$$

a ratio of roughly fourteen to one once each is normalised by its plausible range. The identification is therefore well posed with respect to mass and effectively blind to stiffness. Repeating it across a stiffness range of 40 MN/m to 90 MN/m returns identified masses between 7696 and 7730 t, a spread of 34 t, or 0.4 % of the payload.

A.2.4 Effect on the validation results

The period is used to set the mass, so agreement in period is not independent evidence. The relevant test is whether the identified value improves quantities that played no part in the identification. Table A.2 compares the two candidate masses across the full validation set of Chapter 4.

Table A.2: Validation metrics under the payload mass and the identified mass. None of these quantities entered the identification.

Metric	M_{kin} = 7650 t	M_{kin} = 7719 t
Spectral RMS vs ORCAFLEX, mean abs. error	12.5 %	**10.3 %**
Hinge-driven time domain, mean abs. RMS error	11.7 %	**9.5 %**
Hinge-driven time domain, mean R^2	0.979	**0.986**
Three-hour maximum vs measured record	7.3 %	**5.0 %**
Sea states where the identified mass is closer	—	**6 of 6**

The identified mass is closer on every metric and in every sea state, including the comparison against the measured vessel record, which involves no ORCAFLEX reference at all. Since fitting a period carries no guarantee that response amplitudes, phase correlation or three-hour extremes will also improve, the consistent improvement across these independent quantities supports the interpretation of 7719 t as a physically meaningful participating inertia rather than a curve fit.

A.2.5 Sensitivity of the sway period to the assumed mass

Table A.3 places the choice in proportion. Across the entire plausible swept range of participating mass (7600 t to 7850 t), the sway period varies by only 0.25 s. For scale, this entire span is more than an order of magnitude smaller than the 2.69 s detuning achieved by the hoist optimisation in Chapter 5. No conclusion of this thesis is sensitive to the assumed mass value.

Table A.3: Sensitivity of the sway natural period to the assumed inertial mass. The marginal shift confirms that the sway period is governed by pendulum geometry rather than the exact participating rigging mass.

Mass Assumption	M_{kin} [t]	T_{sway} [s]
Payload mass only	7650	15.63
Identified mass (period-matched)	7719	15.70
Sensitivity gradient		0.10 s per 100 t
Maximum period span over 7600 t to 7850 t sweep		0.25 s

A.3 Insensitivity of the sway mode to hoist-wire stiffness

The hoist lines comprise a steel MH segment and a Dyneema segment acting structurally in series, so their compliances add:

$$\frac{1}{k_{\mathrm{eq}}} = \frac{L_{\mathrm{MH}}}{(EA)_{\mathrm{MH}}} + \frac{L_{\mathrm{Dyn}}}{(EA)_{\mathrm{Dyn}}}. \tag{A.2}$$

For the baseline arrangement, $L_{\mathrm{MH}} = 38.61\,\mathrm{m}$ and $L_{\mathrm{Dyn}} = 23.60\,\mathrm{m}$ give $k_{\mathrm{eq}} = 63\,270\,\mathrm{kN/m}$. The compliance splits 25 % to the steel segment and 75 % to the Dyneema, so the shorter Dyneema segment governs the axial stiffness—its modulus is roughly an order of magnitude lower than steel.

Sweeping the arrangement across the full feasible stiffness range, 37 MN/m to 146 MN/m, gives the modal changes in Table A.4.

Table A.4: Natural period boundaries across the full feasible range of hoist-wire stiffness, 37 MN/m to 146 MN/m.

Mode	Period Span across k Sweep ($T_{n,37\text{MN/m}} \rightarrow T_{n,146\text{MN/m}}$)
Sway	**15.75 s → 15.70 s**
Yaw	4.64 s → 2.42 s
Roll	1.58 s → 0.81 s
Heave	1.51 s → 0.76 s

The sway mode is immune because its restoring action is gravitational rather than elastic. The sway stiffness is $K_{\mathrm{sway}} = 1236\,\mathrm{kN/m}$ against an axial stiffness of 63 270 kN/m—a ratio of roughly fifty to one—so for the pendular mode the wires behave as effectively inextensible. The remaining three modes are genuinely elastic and track the stiffness directly, halving in period across the same range.

Two consequences follow. First, no feasible change to the hoist-wire arrangement—altering the steel-to-Dyneema proportion at constant total length, or shortening the suspension—can detune the sway resonance out of the swell band. Increasing the Dyneema fraction lowers k_{eq} and shortening the line raises it approximately as $1/L$, but the sway period does not respond to either. The sway mode can only be detuned geometrically, through the drop height and the top spread, which is why those are the design variables adopted in Chapter 5.

Second, the provenance of k matters for everything except sway. The heave, roll and yaw periods, and consequently the predicted wire tensions, depend on it directly. The value used here is the theoretical series stiffness of Equation (A.2); it has not been closed against a measured heave decay or static pull test, and that remains a recommended verification.

Appendix B

Free-Decay Validation Details

This appendix documents the free-decay tests used to identify the baseline structural damping and to isolate the contribution of the aerodynamic drag, the winch, the link, and the Block Connection System (BCS). The main text of Chapter 4 only reports the final identified damping ratio and the resulting validation outcome, while the detailed interpretation of the decay envelopes is collected here for completeness.

B.1 OrcaFlex Model Components

In the ORCAFLEX model, the *link* represents the rigid or semi-rigid connection between the lifting arrangement and the suspended payload, while the *winch* represents the tension-controlling element that provides payout and recovery of the hoist line. In the present model, these components affect the dynamic response primarily through the line geometry, the effective constraint on relative motion, and the redistribution of tension during sway motion. The BCS adds additional contact and dissipation effects that are absent in the bare structural configuration.

B.2 Free-Decay Test Matrix

Two sets of free-decay tests were performed. The first set isolated the influence of aerodynamic drag by comparing the nominal case against cases with mean wind speed varied from 1 m/s to 10 m/s and a no-aerodynamic-drag case. The second set isolated the influence of the hoist components by comparing the nominal configuration against cases with the link removed, the winch removed, and the BCS removed. The extracted damping ratios and decay periods are summarised in Table B.1.

Table B.1: Free-decay summary for aerodynamic and hoist-component sensitivity tests.

Config	ζ [%]	T_d [s]	A_0 [m]
No Aero	0.965	15.41	0.54
Nominal (DT1)	0.978	15.41	0.54
Wind 1 m/s	1.035	15.41	0.54
Wind 2 m/s	1.104	15.41	0.54
Wind 10 m/s	1.657	15.41	0.52
No Link	0.839	15.18	0.49
No Winch	0.949	15.62	0.56
No BCS (Bare)	0.514	15.71	0.57
DT2	1.023	15.41	1.62
DT3	1.075	15.41	2.57
DT4 (large)	1.130	15.41	4.11

B.3 Aerodynamic Drag

The aerodynamic-drag study shows that mean wind has only a modest influence on the decay rate of the sway mode. The identified damping ratio increases from 0.965% in the no-aero case to 1.657% at 10 m/s, while the decay period remains unchanged at 15.41 s. This confirms that aerodynamic drag contributes dissipation, but not enough to alter the dominant resonance structure of the suspended jacket in a large way. For the purposes of the reduced-order control plant, this effect is therefore conservative to neglect in the linear model.

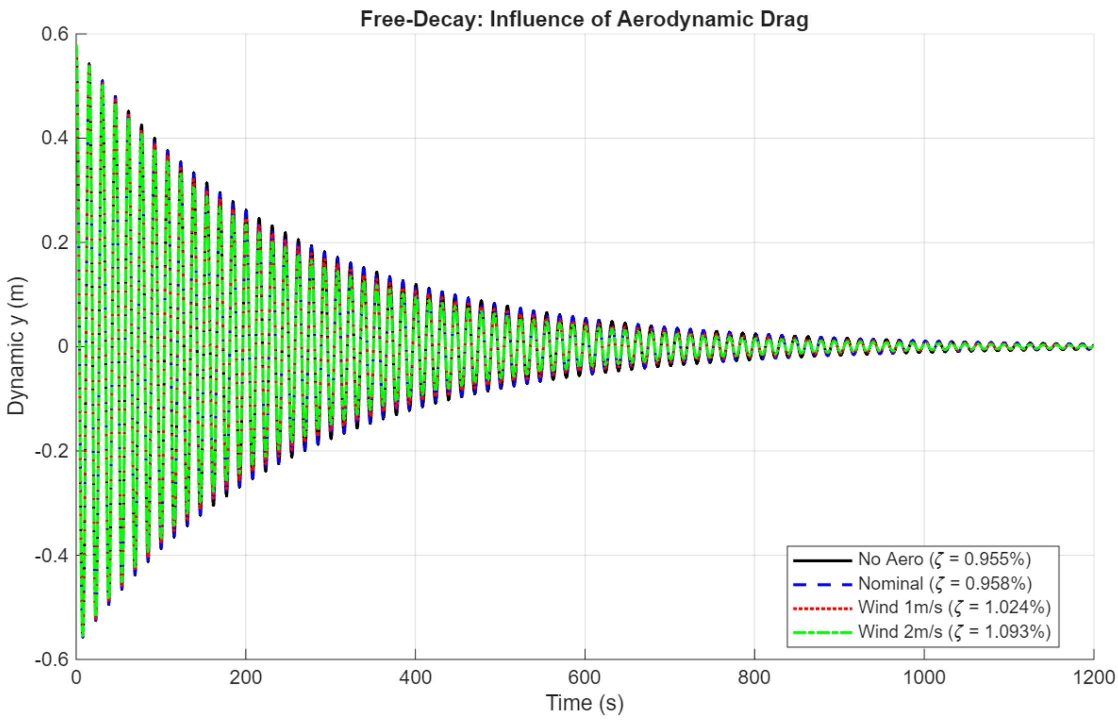

Figure B.1: Free-decay response with aerodynamic drag sensitivity.

B.4 BCS, Winch, and Link Effects

The second test set isolates the contribution of the hoist components. Removing the link reduces the damping ratio to 0.839%, removing the winch gives 0.949%, and removing the BCS gives 0.514%, which is the lowest damping observed in the test matrix This indicates that the BCS is the dominant source of mechanical dissipation among the components tested, while the link and winch each provide a smaller but still measurable contribution to the overall decay behaviour.

The nominal configuration remains close to the no-aero case in terms of damping magnitude, which shows that the baseline sway decay is governed primarily by structural and rigging mechanics rather than by wind. The slight difference between the nominal and no-aero cases is consistent with weak aerodynamic damping under low-velocity oscillation. The no-BCS case is significantly less damped, and therefore exhibits slower energy loss and a longer persistence of oscillation

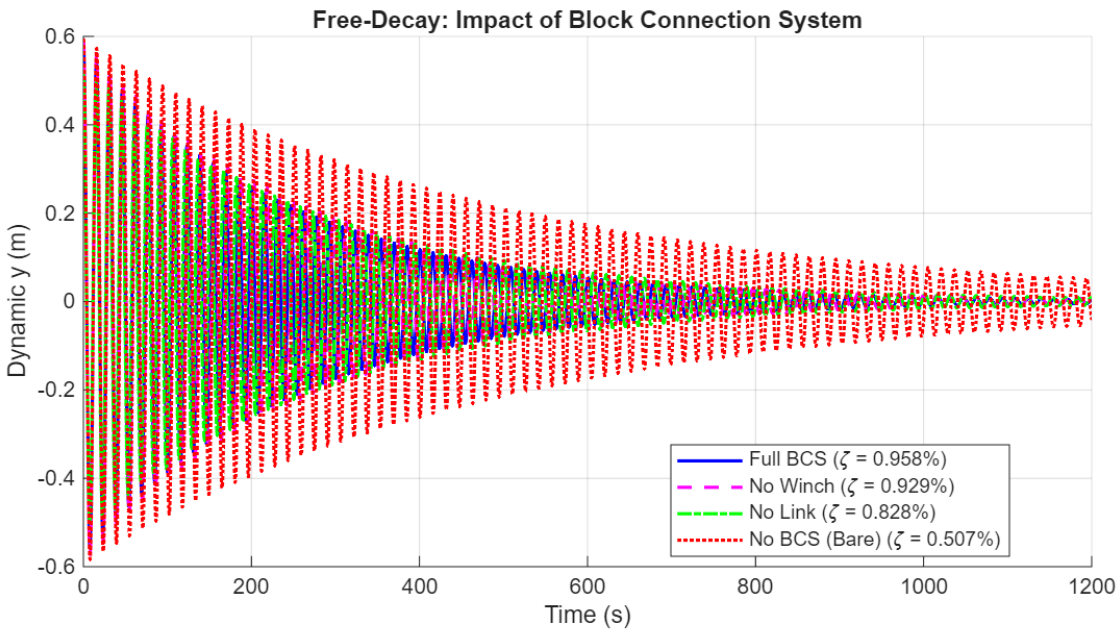

Figure B.2: Free-decay response isolating the BCS, winch, and link contributions.

B.5 Interpretation for Design Sea States

Although the BCS increases dissipation in free-decay, its effect in forced-response sea states is not uniformly beneficial. Across the design sea states used in the thesis, the configuration with BCS enabled generally remains comparable to the cases without BCS and without the winch or link removed, which supports the conclusion that the linear ROM may use a conservative baseline damping representation for control studies However, at higher-frequency periods the BCS can increase the RMS response, indicating that its added mechanical constraints may also amplify certain motion components under wave excitation. For this reason, the BCS is treated as a configuration-dependent effect in the nonlinear verification, but not as a dominant source of damping in the reduced-order control plant

B.6 Implication for the Reduced-Order Model

The free-decay analysis justifies the use of a single effective structural damping ratio in the ROM. The decay tests show that the nominal system is close to 1% damping, with small variations caused by drag and hoist-component configuration, while the fundamental sway period remains stable at approximately 15.41–15.71 s depending on the configuration. This supports the modelling choice used in the main chapters: the ROM retains the dominant sway dynamics while treating the detailed BCS, winch, and link mechanics as higher-order effects that are only reintroduced in the ORCAFLEX verification stage.

Appendix C

Discretisation Convergence

This appendix quantifies the effect of finite wave-component discretisation on the predicted sway response, supporting the discussion in Chapter 4.

Figure C.1 illustrates how the relative sway root-mean-square (RMS) converges as the number of wave components (N) increases for both the V4 and M1 sea states in the baseline ORCAFLEX model (with BCS engaged). The response stabilizes at higher component counts, demonstrating that the discrepancies at longer periods observed in Chapter 4 contain a resolution effect rather than purely a modelling error.

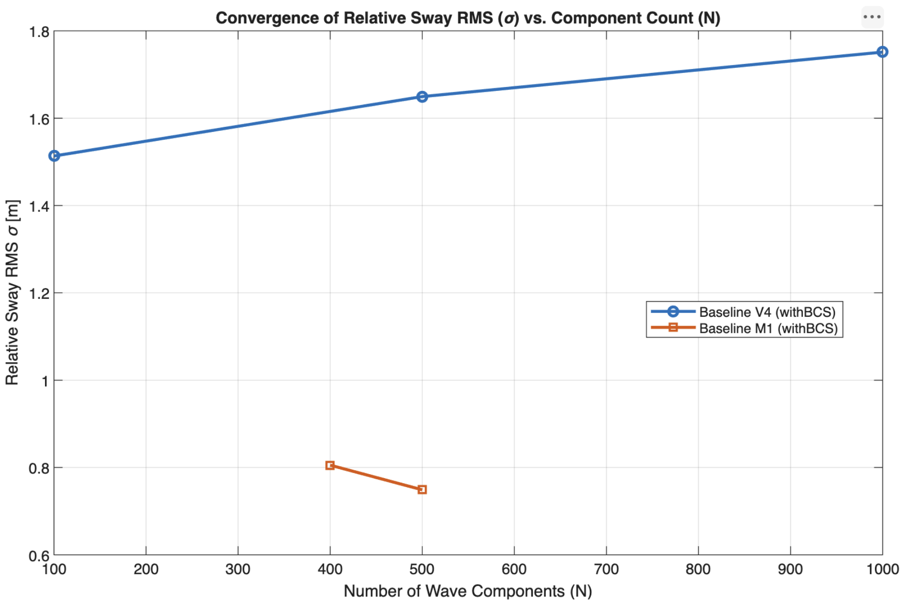

Figure C.1: Convergence of Relative Sway RMS (σ) against the number of wave components (N) for the baseline V4 and M1 sea states (with BCS engaged).