



Design and Optimisation of an Internal Wave Energy Harvester for a Self-Powered Monitoring Buoy

by

Hafþór Árni Hermannsson

Thesis of 30 ECTS credits submitted to the Department of Engineering
at Reykjavík University in partial fulfillment
of the requirements for the degree of
Master of Science (M.Sc.) in Mechatronics Engineering

June 2026

Supervisor:

Baldur Þorgilsson

Examiner:

Indriði Sævar Ríkharðsson

Copyright

Hafþór Árni Hermannsson

June 2026

Design and Optimisation of an Internal Wave Energy Harvester

for a Self-Powered Monitoring Buoy

Hafþór Árni Hermannsson

June 2026

Abstract

Monitoring marine ecosystems and maritime activity in the North Atlantic requires autonomous sensing platforms that can operate for long periods with limited maintenance. Conventional power sources like batteries have finite lifetimes, while solar panels and externally exposed generators are limited by winter darkness, salt buildup, corrosion, bio-fouling, and mechanical damage. This thesis investigates a sealed internal wave-energy harvester that converts buoy motion into electrical power while keeping the main energy-conversion mechanism protected inside the buoy. The selected architecture is a linear permanent-magnet electromagnetic harvester in which a coil assembly translates along a segmented magnetic rod.

The concept is developed through electromagnetic modelling, prototype construction, experimental characterisation, coupled electromechanical validation, and forward scaling analysis. A bench-scale prototype was built with a three-coil carriage and onboard instrumentation for measuring electrical power and carriage motion. A custom rope-driven wave simulator was then developed to provide repeatable vertical excitation for controlled testing. Prototype testing and coupled electromechanical modelling were used to identify mechanical losses, electromagnetic damping, resonance behaviour, and generated power, providing a calibrated basis for evaluating the architecture under larger-scale design assumptions.

Scaled-design simulations based on the validated model then examined whether the architecture could approach the long-term target of 10 W average electrical power for a compact, man-portable monitoring buoy. The simulations used representative sea states from a quality-filtered ocean wave record, with buoy excitation approximated using a simplified wave-follower assumption. The model-based results indicate that reaching the target is technically plausible but demanding, requiring increased moving mass, improved electromagnetic coupling, careful stroke management, and load-tuning strategies. Within a strict 1.8 m rod-length envelope, the best scaled candidate predicted an output of about 8.8 W under the screening assumptions, while larger packages reached the 10 W class. The results support further development of the selected architecture, with the main remaining challenges being a practical stroke-management mechanism, mechanical integration, and further buoy-coupled hydrodynamic modelling.

Keywords: Wave energy harvesting, autonomous monitoring buoy, permanent-magnet linear generator, coupled electromechanical modelling, experimental characterisation, scaled-design analysis.

Contents

List of Figures	viii
List of Tables	xii
List of Abbreviations	xiv
List of Symbols	xvi
1 Introduction	1
1.1 Background and Motivation	1
1.2 Problem Statement	2
1.3 Research Objectives	2
1.4 Scope and Contributions	3
1.5 Thesis Structure	3
2 State of the Art	5
2.1 Scope of the Literature Review	5
2.2 Power Supply for Ocean Monitoring Systems	5
2.3 Sealed and Enclosed Wave Energy Conversion	6
2.4 Electromagnetic Harvesters for Buoys and Internal WECs	6
2.5 Alternative Transduction Technologies	7
2.6 Low-Frequency and Irregular-Wave Motion Conditioning	8
2.7 Generator Design, Sealing, and Electromagnetic Detail	8
2.8 Modelling and Experimental Validation Approaches	9
2.9 Local Icelandic and Reykjavík University Context	9
2.10 Research Gap and Thesis Scope	10
3 Methodology and Design Concept	11
3.1 Design Process Overview	11
3.2 Concept Selection	11
3.2.1 Scope of Considered Concepts	11
3.2.2 Architecture Families Considered	11
3.2.3 Cross-Cutting Dynamic Strategies in the Literature	12
3.2.4 Representative Reference Concepts from the Literature	12
3.2.5 Evaluation Criteria and Decision Matrix	13
3.2.6 Selection Outcome	14
3.3 Selected Architecture	15
4 Electromagnetic Modelling and Coil Optimisation	17
4.1 Modelling Scope and Strategy	17

4.2	Magnetic Rod Characterisation as Model Input	18
4.2.1	Rod Identification and External Geometry	18
4.2.2	Surface Magnetic Field Measurements	18
4.2.3	Field-Derived Rod Geometry	19
4.2.4	Analytical Interpretation of the Measurements	20
4.3	Finite-Element Field Reconstruction	21
4.3.1	ANSYS Maxwell Model	21
4.3.2	MATLAB PDE Toolbox Model	22
4.4	Field-Map Calibration and Agreement	23
4.5	Coil-Level Electromagnetic Calculations	24
4.5.1	Flux Linkage from the Vector Potential	25
4.5.2	Induced Voltage and Matched-Load Power	25
4.6	Coil Outer-Diameter Optimisation	25
4.6.1	Cross-Validation Result	26
4.6.2	Sensitivity to Winding Fill Factor	27
4.6.3	Recommended Coil Outer Diameter	28
4.7	Adopted Electromagnetic Model	28
4.8	Chapter Summary	28
5	Prototype Development and Experimental Platform	31
5.1	Role of the Prototype in the Thesis	31
5.2	Mechanical Implementation of the Selected Architecture	31
5.2.1	Support Structure, Carriage, and Moving Mass	32
5.2.2	Prototype Coil Geometry and Assembly	32
5.2.3	Coil Fabrication	34
5.2.4	Elastic Restoring-Rope Arrangement	34
5.3	Measurement and Data-Acquisition System	36
5.3.1	Electrical Measurement Paths	37
5.3.2	Motion Measurement and Translation Estimation	37
5.3.3	Embedded Interface, Sampling, and Logging	38
5.4	Integrated Prototype Assembly	39
5.5	Wave-Simulator System for Controlled Excitation	40
5.5.1	Mechanical Principle and Motion Sensing	41
5.5.2	Control Architecture and Operating Modes	42
5.5.3	Supervision, Safety, and Key Parameters	43
5.6	Coupled Simulator–Harvester Logging	45
5.7	Chapter Summary	46
6	System-Level Dynamic Model	47
6.1	Purpose and Scope	47
6.2	Coordinate Definition and Modelling Assumptions	47
6.3	Nonlinear Restoring Force	48
6.4	Damping Model	48
6.5	Equation of Motion	49
6.6	Linearised Analysis	49
6.6.1	Small-Amplitude Natural Frequency	50
6.6.2	Base-Excitation Transmissibility	50
6.6.3	Coupled Electrical Loading and Useful Power	50
6.7	Nonlinear Frequency Response	51

6.8	Model Predictions to be Tested Experimentally	51
6.8.1	Quantities Predicted from Geometry and Rope Model	51
6.8.2	Predictions Requiring Experimentally Identified Parameters	51
6.8.3	Summary of Testable Predictions	52
6.9	Chapter Summary	53
7	Experimental Setup and Results	55
7.1	Experimental Objectives	55
7.2	General Experimental Methodology	55
7.2.1	General Test Protocol	55
7.2.2	Instrumentation Accuracy and Known Limitations	55
7.3	Free-Decay Experiments	56
7.3.1	Purpose	56
7.3.2	Procedure	57
7.3.2.1	Load-dependent tests	57
7.3.2.2	Multi-amplitude open-circuit tests	57
7.3.3	Analysis Method	57
7.3.4	Results: load-dependent campaign	59
7.3.5	Results: multi-amplitude campaign	61
7.3.6	Discussion	64
7.4	Steady-State Wave-Driven Experiments	65
7.4.1	Purpose and Analysis Strategy	65
7.4.2	Benchmark Resonance-Region Load Sweep	66
7.4.3	Frequency-Response Experiments	67
7.4.3.1	Purpose and Procedure	67
7.4.3.2	Analysis Method	68
7.4.3.3	Results and Discussion	69
7.4.4	Amplitude-Response Experiment	72
7.4.4.1	Purpose and Procedure	72
7.4.4.2	Analysis Method	72
7.4.4.3	Results and Discussion	73
7.5	Chapter Summary	75
8	Coupled Model Validation	77
8.1	Purpose and Scope	77
8.2	Coupled Model Used for Validation	78
8.3	Validation Strategy	78
8.4	Electromagnetic Replay Validation	79
8.4.1	Method	79
8.4.2	Results	79
8.5	Loss-Partition Transfer Test	80
8.5.1	Method	80
8.5.2	Results	80
8.6	Final Fixed-Parameter Coupled Validation	81
8.6.1	Model configuration	81
8.6.2	Overall result	81
8.6.3	Interpretation of the Remaining Discrepancy	84
9	Scaling Analysis and Design Implications	85

9.1	Scaling Workflow and Model Scope	85
9.2	Sinusoidal Screening	87
9.3	Sea-State Representation and Ocean-Weighted Power	89
9.4	Wave-Follower Ocean Screening	91
9.5	Stroke Management	92
9.6	Electrical Load Tuning	96
9.7	Summary of Scaling Results	100
10	Conclusions and Future Work	105
10.1	Overview	105
10.2	Summary of the Work	105
10.3	Main Conclusions	106
10.4	Answers to the Research Questions	108
10.5	Limitations	109
10.6	Recommended Future Work	111
10.7	Final Remarks	112
	Data and Code Availability	113
	Bibliography	115

List of Figures

4.1	Comparison of two raw surface-field measurement passes over one magnetic repeat period.	19
4.2	CAD visualisation of the inferred segmented magnetic-rod geometry. Red and blue indicate opposite magnetisation directions; grey indicates spacers and end regions.	20
4.3	Analytical comparison of two radial geometry interpretations. The thick-wall interpretation smooths the transition peaks too strongly, while the thin-wall interpretation better reproduces the measured radial-field gradients near the magnet-spacer transitions.	21
4.4	Magnetic field-line plot from the ANSYS Maxwell model of a representative 7-magnet subsection. The alternating bucking magnetisation concentrates flux near the magnet-spacer transition regions.	22
4.5	Radial magnetic field component $B_r(r, z)$ over the full 1010 mm rod from the MATLAB PDE Toolbox model. The periodic bucking structure is visible along the rod, with reduced field amplitude near the end regions.	23
4.6	Radial and axial magnetic field components along a line just outside the rod surface, evaluated over the full rod length. The central region is nearly periodic, while the end regions show the expected reduction in field amplitude.	23
4.7	Comparison between measured and modelled field components over one magnetic repeat period at the surface-probe radius. The radial component shows the paired peak structure used to infer the magnet-spacer geometry, while the axial component provides an additional check of the field variation through the magnet bodies.	24
4.8	Predicted matched-load power $P_{\max}$ as a function of coil outer diameter, comparing the ANSYS Maxwell and MATLAB PDE field maps at $v = 0.5$ m/s and $f = 0.85$. Both models identify the same optimum at $OD \approx 48.7$ mm.	26
4.9	Sensitivity of predicted matched-load power to winding fill factor using the MATLAB PDE field map. The predicted power level increases with fill factor, while the optimum coil outer diameter shifts by less than 0.5 mm across the tested range.	27
5.1	Bench-scale prototype assembly showing the segmented magnetic rod and translating three-coil carriage.	31
5.2	Three-coil carriage geometry and fabricated prototype assembly.	33
5.3	Custom coil-winding rig used to manufacture the prototype coils.	34
5.4	Unified empirical tension-strain model for the elastic ropes, obtained by pooling cyclic loading and unloading data from the three tested free lengths.	35
5.5	Measurement and data-acquisition circuit used for the bench-scale prototype.	36
5.6	Fabricated prototype measurement and data-acquisition electronics used during bench testing.	38

5.7	Fully assembled moving carriage, showing the attached masses and onboard measurement electronics.	39
5.8	Top-down view of the fully integrated bench-scale prototype, showing the segmented magnetic rod, translating three-coil carriage, onboard measurement electronics, attached masses, and elastic restoring ropes.	39
5.9	Wave-simulator rig used to impose controlled vertical excitation during laboratory testing of the harvester prototype. The motor-driven wheel converts rotary motion into vertical rope displacement, while the counterbalance reduces the one-sided gravitational bias from the suspended load.	40
5.10	Wheel configuration used in the wave simulator.	41
5.11	Wave simulator circuit schematic, showing the Raspberry Pi Zero 2 W controller, motor-driver interface, 24 V motor power path, and wheel-motion sensing hardware.	42
5.12	High-level software and signal-flow architecture of the wave simulator, showing the browser dashboard, Raspberry Pi backend, motor-control path, wheel sensing, safety supervision, experiment logging, and generator-monitor communication.	43
5.13	Browser dashboard used for coupled simulator–harvester experiments. The left side controls the wave simulator, experiment logging, frequency sweeps, and safety limits, while the right side displays the ESP32-based generator monitor with live electrical measurements, load settings, power estimates, and carriage translation.	45
7.1	Load-dependent free-decay response. The displacement traces show the effect of electrical loading on the decay rate, while the damping plot shows the loaded cases following the expected inverse-resistance trend.	59
7.2	Free-decay damping and bypass-path energy. The inferred electromagnetic damping increases as load resistance is reduced, while the bypass-path energy is highest at an intermediate load.	60
7.3	Multi-amplitude open-circuit decay summary. Top: release repeatability and fitted natural frequency. Bottom: open-circuit damping ratio and voltage–speed estimate of K_{emf} . Marginal groups are plotted with hollow markers and excluded from the regression.	62
7.4	Group-level Coulomb-plus-viscous decomposition of the open-circuit damping. (a) Equivalent damping coefficient versus release amplitude with the fitted Coulomb–viscous model from Equation (7.11). (b) The same data plotted against $1/X$, where the Coulomb component appears as a linear trend.	63
7.5	Cycle-resolved Coulomb-plus-viscous decomposition. Each marker represents one same-sign extrema pair from one decay run, with X_{pair} taken as the geometric mean of the two peak amplitudes.	63
7.6	Benchmark resonance-region load-sweep results. (a) Mean measured bypass-path power as a function of external bypass-path resistance. (b) Measured simulator amplitude and fitted carriage translation amplitude as functions of external bypass-path resistance. Error bars show the standard deviation across repeated runs.	67

7.7	Internal harvester frequency response in bypass mode at $R_L = 9.8 \Omega$. (a) Fundamental transmissibility X_1/A_1 versus forcing frequency for both campaigns, compared with the SDOF prediction from the decay-identified parameters and an independent SDOF fit to the variable-amplitude data. (b) Mean bypass-path power versus forcing frequency.	69
7.8	Constant-amplitude sweep summary at $A \approx 0.4$ m. (a) Achieved simulator amplitude versus period. (b) Mean bypass-path power versus period. (c) Carriage half peak-to-peak amplitude versus period. (d) Transmissibility from both the FFT fundamental and the peak-to-peak amplitude, compared with the SDOF prediction at the decay-identified parameters. The separation between the two transmissibility curves indicates harmonic contribution to the carriage motion.	70
7.9	Harmonic-resonance mechanism in the constant-amplitude sweep. (a) Carriage harmonic content from 1f to 5f by forcing period. (b) Harmonic frequencies relative to the decay-identified natural frequency. Marker size is proportional to the carriage amplitude at that harmonic.	71
7.10	Measured bypass-path power versus reconstructed RMS carriage velocity squared across the frequency-sweep campaigns. The near-linear collapse supports the linear electromagnetic damping assumption used in Chapter 6.	72
7.11	Reconstructed carriage motion over one normalised forcing cycle for each tested commanded amplitude at $T = 4$ s.	73
7.12	Mean bypass-path power versus achieved base amplitude at $T = 4$ s and $R_L = 9.8 \Omega$. The fitted power law is compared with a linear-system $P \propto A^2$ reference. The shaded region marks the simulator-tracking-limited part of the sweep.	74
8.1	Workflow used for the coupled-model validation and subsequent scaling analysis. The replay step evaluates the electromagnetic subsystem using measured carriage motion. The coupled validation then predicts carriage motion from measured base excitation, and the fixed harvester-side model is carried into the forward scaling analysis.	77
8.2	Electromagnetic replay validation. Measured carriage motion is prescribed as the input to the PDE-derived electrical model.	79
8.3	Final fixed-parameter coupled validation. The amplitude comparison evaluates the predicted carriage stroke, while the power comparison includes the combined effect of predicted motion, velocity, magnetic phase, and electrical loading.	82
8.4	Representative time-domain coupled-validation result for the A0.4_T3.5 case at $R_L = 9.8 \Omega$. The measured simulator base motion is shown together with the measured carriage displacement and the coupled-model prediction.	82
8.5	Representative measured and predicted instantaneous generator-voltage snapshot for the same A0.4_T3.5 validation case. The model signal is sampled in the same manner as the logged <code>vgen_inst</code> channel.	83
9.1	S1 and S2 screening under the 4 s sinusoid. (a) S1 design-space map. (b) S2 power versus coil count at $f_n = 0.25$ Hz, grouped by moving mass.	88
9.2	S2 and S3 long-rerun results in the $(L_{rod}, \bar{P}_L)$ plane. The highest-power sinusoidal point is the S2 4 s rerun with $m_{EM} = 1.5$. Several other S2 candidates remain near 9 W to 10 W, but the longer-period S3 cases stay below 5 W.	89
9.3	Quality-filtered Surtseyjardufl (H_s, T_z) distribution and the ten representative states. Marker area scales with normalised occurrence weight.	90

9.4	O1 wave-follower ocean screen: ocean-weighted external-load power against worst-case required rod length. The strongest candidates exceed the 1.8 m envelope.	92
9.5	O2 feasibility-first ocean screen. (a) Raw and 1.8 m-cutout power per final candidate. (b) Power against worst-case required rod length.	92
9.6	B1 stroke-management diagnostic under cylinder-RAO base motion. (a) Cubic hardening. (b) Soft-stop bumpers. Harder force laws reduce package length and recovered power.	95
9.7	B2 integrated bumper sweep. (a) Ocean-weighted power against peak relative carriage acceleration. (b) Best B2 candidate per package class.	96
9.8	Load-tuning recovery from B2 to L1. The LR00001 bar uses the confirmed ten-state strict-envelope result.	98
9.9	L1 feasibility landscape. Ocean-weighted external-load power is plotted against worst-case required rod length. Strict-1.8 m confirmation rows are overlaid. . .	99
9.10	Per-sea-state weighted contribution $P_{\text{mean}}\omega$ for three L1 reference rows: (a) LR00241, (b) LR00169, and (c) strict-envelope LR00001. Operational states are SS01–SS06 and stress states are SS07–SS10.	99
9.11	Combined feasibility map across all scaling stages. Ocean-weighted external-load power is plotted against worst-case required rod length. The shaded target region is $L_{\text{rod}} \leq 1.8 \text{ m}$ and $\bar{P}_L \geq 10 \text{ W}$. No candidate lies inside the strict target region.	101
9.12	Explored design space for the 251 unique candidates included in the combined Pareto summary. Each line represents one candidate. The axes show the screening stage, main design parameters, stroke-management settings, and binned outcomes for power and required rod length. Line colour indicates predicted mean external-load power.	102

List of Tables

3.1	Evaluation criteria used to compare candidate internal architecture families.	13
3.2	Weighted decision matrix for the candidate internal architecture families. Scores are on a 1–5 scale; criteria C1–C5 correspond to those in Table 3.1 in the same order.	14
4.1	Segmented magnetic-rod geometry used in the model.	20
4.2	Field-level calibration summary against measured Hall-probe data over one magnetic repeat period.	24
4.3	Coil OD optimisation summary for the two FEM field maps, evaluated at $v = 0.5$ m/s and $f = 0.85$	27
5.1	Main mechanical parameters of the bench-scale prototype.	32
5.2	Main coil parameters of the bench-scale prototype.	33
5.3	Main components of the prototype monitoring electronics.	36
5.4	Selected implemented parameters of the wave simulator.	44
6.1	Summary of model predictions and the experimental measurements used to evaluate them.	52
7.1	Electrical loads used in the load-dependent free-decay campaign.	57
7.2	Summary of extracted free-decay metrics for each electrical load. Reported values are mean $\pm$ standard deviation over six runs.	60
7.3	Damping decomposition from the load-dependent free-decay tests. The listed R_{tot} includes the measured parallel coil-bank resistance and the external circuit resistance, and c_{em} is obtained by subtracting the open-circuit mechanical baseline from the loaded-case total damping.	61
7.4	Multi-amplitude open-circuit free-decay summary. Values are mean $\pm$ standard deviation over four runs per amplitude group. The 5 cm group is marginal because its decay records contain three or fewer extrema.	62
7.5	Sensitivity of the cycle-resolved decomposition to the amplitude-pair mean. The geometric-mean variant is adopted as the preferred result; the arithmetic and harmonic variants are retained as bounds.	64
7.6	Key parameters identified from the free-decay experiments.	65
7.7	Benchmark resonance-region load-sweep summary. Values are mean $\pm$ standard deviation across three repeated runs.	66
7.8	Amplitude-sweep summary at $T = 4$ s and $R_L = 9.8 \Omega$. The simulator delivery is the ratio of achieved to commanded base amplitude. The effective natural frequency is inferred from the fundamental transmissibility under a linear SDOF assumption.	75

8.1	Electromagnetic replay validation metrics. One global scale factor is fitted on the training subset and applied unchanged to the holdout subset.	80
8.2	Loss-partition transfer test on the full 52-run wave-driven dataset.	80
8.3	Fixed parameter set used in the final coupled validation.	81
8.4	Overall final coupled-validation metrics across the full 52-run wave-driven dataset.	83
8.5	Final coupled-validation results by experimental group.	83
9.1	Summary of the scaling stages used in this chapter.	87
9.2	Headline sinusoidal screening results. Power values are rounded to one decimal.	88
9.3	Representative sea states from the Surtseyjardufl record. Weights are normalised over all ten states.	89
9.4	Headline O1 and O2 results. Power values are rounded to one decimal.	91
9.5	Stroke-management results. Values are rounded summary results. B1 uses cylinder-RAO base motion; B2 uses wave-follower base motion with integrated bumper dynamics.	94
9.6	Bumper settings for the headline controlled-stroke candidates.	95
9.7	Headline L1 results. The $\pm$ value is the seed-to-seed standard deviation of the ocean-weighted total.	97
9.8	Best 1.8 m envelope results across the campaign, all on the full ten-state weighting.	103

List of Abbreviations

Abbreviation	Meaning
ADC	Analog-to-digital converter
ANSYS	Ansys Maxwell finite-element software
CAD	Computer-aided design
CSV	Comma-separated values
DC	Direct current
EM	Electromagnetic
EMF	Electromotive force
FEM	Finite element method
FFT	Fast Fourier transform
HTTP	Hypertext Transfer Protocol
I ² C	Inter-Integrated Circuit
ID	Identifier
IMU	Inertial measurement unit
ISWEC	Inertial Sea Wave Energy Converter
MATLAB	Matrix Laboratory software
OD	Outer diameter
ODE	Ordinary differential equation
PDE	Partial differential equation
PID	Proportional–integral–derivative
PTO	Power take-off
PTFE	Polytetrafluoroethylene
PWM	Pulse-width modulation
RAO	Response amplitude operator
RMSE	Root-mean-square error
RMS	Root mean square
SDOF	Single degree of freedom
TFT	Thin-film transistor
THD	Total harmonic distortion
WEC	Wave energy converter

List of Symbols

Symbol	Description	Units
A	Imposed, measured, or fitted displacement amplitude	m
A_ϕ	Azimuthal magnetic vector potential	T m
a_{peak}	Peak carriage acceleration	multiples of g
B_{rem}	Remanent flux density	T
B_r, B_z	Radial and axial magnetic flux-density components	T
c_{em}	Equivalent electromagnetic damping coefficient	N s/m
$c_{\text{eq}}, c_{\text{tot}}$	Equivalent and total damping coefficients	N s/m
c_{visc}	Viscous damping coefficient	N s/m
Δf	Frequency spacing in irregular-wave synthesis	Hz
$\Delta z_{\text{phase, sensor}}$	Sensor-frame phase offset used in EM replay	m
δ, δ_k	Logarithmic decrement and cycle-resolved logarithmic decrement	–
$\eta(t)$	Sea-surface displacement used as wave-follower base motion	m
f, f_k	Frequency and spectral component frequency	Hz
$f_n, f_{n, \text{target}}$	Natural frequency and target natural frequency	Hz
F_c	Coulomb friction force	N
F_{em}	Electromagnetic reaction force	N
F_{rest}	Restoring force	N
F_{rope}	Rope restoring force	N
F_{stop}	Modelled end-stop force	N
g	Gravitational acceleration	m/s ²
H_s	Significant wave height	m
I, I_{shunt}	Current and shunt current	A
$K(q)$	Position-dependent electromagnetic coupling coefficient	V s/m
K_{emf}	Adopted voltage–speed constant	V s/m
k	Power–velocity coefficient in $P = kv_{\text{rms}}^2$	W s ² /m ²
k_3	Cubic restoring stiffness coefficient	N/m ³
k_{eff}	Effective linear restoring stiffness	N/m
L_{coil}	Coil active length	m
L_{rod}	Required rod length	m
λ	Flux linkage	V s
m	Moving carriage mass	kg
m_{EM}	Electromagnetic coupling multiplier	–
μ_0, μ_r	Vacuum permeability and relative permeability	N/A ² , –
N, N_f, N_{turns}	Coil count, number of spectral components, and number of coil turns	–
$\omega, \omega_d, \omega_n$	Forcing, damped natural, and natural angular frequency	rad/s
$\bar{P}_L$	Mean external-load power	W
$\bar{P}_{L, \text{ocean}}$	Ocean-weighted mean external-load power	W
P, P_{load}	Load power	W
P_{max}	Matched-load power	W
p_{bypass}	Bypass-path power	W
$p_{\text{mag}}, p_{\text{struct}}$	Magnetic and structural pitch of the rod	m

Φ	Magnetic flux through one turn	V s
ϕ_k	Random phase of the k -th spectral component	rad
q	Relative carriage displacement	m
$\dot{q}$	Relative carriage velocity	m/s
$\ddot{q}$	Relative carriage acceleration	m/s ²
q_{stop}	Stop position	m
r	Frequency ratio	–
$R_{\text{coil}}, R_{\text{coil,par}}, R_{\text{coil,branch}}$	Coil, measured parallel coil-bank, and branch resistance	Ω
R_L	External load resistance	Ω
R_{shunt}	Current-sense shunt resistance	Ω
R_{circuit}	External circuit or bypass-path resistance, excluding coil resistance	Ω
R_{tot}	Total current-loop resistance including coil, load, and shunt resistance	Ω
$S(f)$	One-sided wave displacement spectrum	m ² /Hz
T	Wave, simulator, or imposed-motion period	s
T_p	Peak wave period	s
T_z	Zero-upcrossing wave period	s
$V, V_{\text{gen}}, V_{\text{rect}}, V_{\text{shunt}}$	Voltage, generator voltage, rectifier voltage, and shunt voltage	V
v, v_{rms}	Velocity and RMS velocity	m/s
v_s	Smoothing velocity in the Coulomb-friction term	m/s
y_{base}	Prescribed base displacement	m
$\dot{y}_{\text{base}}$	Base velocity	m/s
$\ddot{y}_{\text{base}}$	Base acceleration	m/s ²
z, z_0	Axial coordinate and coil offset in the field map	m
$\zeta, \zeta_{\text{eff}}$	Damping ratio and effective damping ratio	–

Chapter 1

Introduction

1.1 Background and Motivation

Autonomous ocean-monitoring systems are increasingly important for marine science, environmental monitoring, fisheries management, and maritime security. In the North Atlantic, and particularly around Iceland, such systems must operate in a harsh and highly variable environment where maintenance visits can be expensive, weather-limited, and disruptive to long-term data collection. For a monitoring buoy, the power system is therefore a central design constraint, affecting mission endurance, operational cost, and practical usefulness.

Conventional power options each have significant limitations in this setting. Primary batteries provide simple and reliable energy storage, but their finite capacity limits the mission duration unless the buoy is recovered or serviced. Solar panels can extend endurance, but their output is strongly affected by day–night cycling, winter darkness, cloud cover, icing, salt deposits, and the orientation and exposure of the buoy. Small wind generators can also be considered, but exposed rotating machinery introduces additional vulnerability to corrosion, fouling, impact, and mechanical wear. These issues are especially relevant for small or man-portable buoys, where available mass, deck area, and maintenance margin are limited.

Ocean-wave energy is attractive because the buoy is already forced by the same environment that it is intended to monitor. The challenge is that the useful motion is irregular, low-frequency, multi-directional, and strongly dependent on the buoy dynamics. Large grid-connected wave energy converters are designed around the extraction of significant wave power, but the present application is different: the objective is to provide a local, robust, low-maintenance power source for an autonomous buoy. For this reason, the most relevant technologies are not necessarily the highest-power wave energy converters, but compact harvesters that can convert buoy motion into electrical power without compromising survivability.

A sealed internal architecture is a promising route toward this objective. In such a system, the main power take-off mechanism is enclosed inside the buoy hull, so the generator, mechanical transmission, bearings, springs, and electrical components are protected from seawater exposure, biofouling, and direct impact. The penalty is that the harvester can no longer interact directly with the wave surface; it must instead extract energy from relative motion between the buoy structure and an internal inertial or compliant subsystem. This changes the design problem from one of external wave loading to one of internal dynamic coupling, electromagnetic conversion, and mechanical packaging.

This thesis investigates a sealed internal wave energy harvester intended for a self-powered monitoring buoy. The selected concept is a linear electromagnetic harvester in which a translating coil carriage moves relative to a segmented magnetic rod. The long-term application target is a self-powered, man-portable monitoring buoy with an average power level on the order of 10 W. The present work does not claim to deliver a final deployed buoy system. Instead, it develops and tests a bench-scale generator architecture, establishes an electromagnetic and dynamic modelling chain, and uses the validated results to assess what design changes would be required for future scaling.

1.2 Problem Statement

The engineering problem addressed in this thesis is the design, modelling, and experimental evaluation of a wave-driven kinetic energy harvester that can operate entirely inside a sealed buoy enclosure. The harvester must convert externally imposed buoy motion into internal relative motion, convert that motion into electrical power, and remain mechanically simple enough to be a credible candidate for long-term use in harsh coastal or offshore conditions.

The central difficulty is the mismatch between the available excitation and the desired power level. Ocean buoys experience slow and irregular motion, while compact electromagnetic generators generally produce more power when relative velocity, magnetic interaction length, and active conductor volume are increased. Many buoy-scale harvesters reported in the literature are aimed at low-power sensing or intermittent self-powered operation [1], [2], [3], [4], while watt-level sealed linear electromagnetic buoy harvesters have been demonstrated more rarely [5]. Moving toward 10 W average power therefore requires more than a small demonstrator: it requires a model that connects wave-driven motion, mechanical dynamics, electromagnetic coupling, electrical loading, and geometric scaling.

The specific problem is therefore to determine whether a sealed internal linear electromagnetic architecture can provide a plausible path from a bench-scale prototype toward a higher-power monitoring-buoy supply. This requires answering both a design question and a validation question. The design question is whether a mechanically simple internal architecture can be selected and sized in a way that is compatible with low-frequency wave excitation. The validation question is whether the resulting model predicts the measured prototype behaviour well enough to support scaling conclusions.

1.3 Research Objectives

The main objective of this thesis is to design, model, build, and experimentally validate an internal kinetic wave energy harvester, and to use the resulting model to assess whether the selected architecture has a technically plausible route toward approximately 10 W average power for a self-powered, man-portable monitoring buoy in Icelandic coastal waters.

The work is guided by the following research questions:

1. **RQ1:** What is the most suitable and cost-efficient internal mechanical architecture for harvesting kinetic energy within a sealed, volume-constrained buoy enclosure?
2. **RQ2:** How can the selected architecture be represented by a coupled electromagnetic and mechanical model, and what predictions does the model make about its behaviour?

3. **RQ3:** How well does the model reproduce the measured electrical and mechanical behaviour of the physical prototype under controlled laboratory excitation?
4. **RQ4:** Based on the validated model, what design changes are required for the concept to move toward the target average power level of approximately 10 W?

1.4 Scope and Contributions

The scope of the thesis is limited to the internal generator concept, its bench-scale implementation, and the modelling and experimental framework required to evaluate it. Full buoy hydrodynamics, survivability testing at sea, long-duration corrosion testing, and final power-electronics integration are outside the present scope. These topics remain important for a deployable system, but the immediate objective is to establish whether the selected internal energy-conversion architecture is worth carrying forward.

The main contributions of the thesis are:

1. the development of a sealed internal linear electromagnetic harvester architecture based on a translating coil carriage and segmented magnetic rod;
2. a reverse-engineered and experimentally calibrated magnetic-field model of the segmented magnetic rod used in the prototype;
3. a coil-level electromagnetic modelling workflow used to select and justify the prototype coil geometry;
4. a bench-scale physical prototype with onboard electrical and motion measurement;
5. a controlled wave-simulator test platform for repeatable excitation of the prototype;
6. experimental characterisation using free-decay and wave-driven tests under different electrical loading conditions;
7. coupled model validation against measured prototype motion and electrical power;
8. a scaling analysis that connects the validated harvester model to the long-term target of a self-powered monitoring buoy.

1.5 Thesis Structure

This thesis is organised as follows. Chapter 2 reviews the relevant state of the art in wave energy harvesting for autonomous buoy systems, with emphasis on sealed or enclosed internal mechanisms, electromagnetic conversion, low-frequency motion conditioning, and local Icelandic buoy-power work. Chapter 3 presents the design methodology and the concept-selection process that led to the selected internal linear sliding architecture. Chapter 4 develops the electromagnetic modelling workflow for the segmented magnetic rod and the coil-level optimisation used to define the prototype generator geometry. Chapter 5 describes the physical prototype, the embedded measurement system, and the wave-simulator platform used for controlled excitation.

Chapter 6 develops the system-level dynamic model by combining the mechanical restoring force model with the electromagnetic damping and electrical loading framework. Chapter 7 presents the experimental setup and prototype results, including free-decay tests,

wave-driven experiments, load-dependent response, and frequency-response measurements. Chapter 8 validates the coupled electromechanical model against the measured prototype response. Chapter 9 uses the validated model to examine the scaling implications and the design changes required to move toward the long-term target of a self-powered monitoring buoy. The final chapter summarises the main conclusions and outlines recommended future work.

Chapter 2

State of the Art

2.1 Scope of the Literature Review

The literature relevant to this thesis spans several overlapping fields: large-scale wave energy converters (WEC), low-power marine energy harvesters, autonomous buoy systems, sealed or enclosed inertial harvesters, electromagnetic generators, nonlinear vibration harvesters, and local Icelandic buoy research. The present thesis is not concerned with grid-scale wave energy conversion in general. Its focus is narrower: compact wave-driven energy harvesting for a self-powered monitoring buoy, with particular emphasis on concepts that can keep the power take-off protected inside a sealed hull.

For that reason, the literature is interpreted according to four questions. First, does the system keep the power take-off mechanism protected inside a hull or sealed enclosure? Second, is it intended for a buoy, ocean sensor, or autonomous marine platform rather than for grid-scale power production? Third, what mechanical architecture and transduction method does it use? Fourth, does it address the main dynamic difficulty of the problem: extracting useful energy from low-frequency, irregular, and often multi-directional wave excitation?

2.2 Power Supply for Ocean Monitoring Systems

The need for autonomous power supply is well established in marine monitoring. Ocean observation buoys and subsea sensors are often deployed to collect environmental, geophysical, biological, or infrastructure-monitoring data over long periods. Battery-only systems are simple, but the requirement for periodic replacement or recharging increases cost and interrupts measurements. Chiu et al. discuss this issue for stationary subsea sensors and motivate energy harvesting as a way to support long-duration monitoring without repeated recovery operations [6]. Cortes et al. demonstrate the same broader need at system scale through the Open Sea Lab observatory in Chile, where wave energy is integrated into coastal sensing infrastructure and the seasonal variability of the harvested power is explicitly reported alongside biofouling growth on mooring lines and junction boxes [7].

The power-supply problem for a monitoring buoy is different from that of a utility-scale WEC. For ocean monitoring, the key requirements are survivability, low maintenance, compatibility with sensors and communications, and sufficient average power for the payload. Depending on payload and duty cycle, the required average power can range from milliwatt- to sub-watt levels for low-power sensing and intermittent telemetry [1], [3], to several watts or more for continuously powered sensors, acoustic systems, data logging,

and communication hardware [5], [6]. Compact wave energy harvesting is therefore attractive, but the device must also satisfy strict constraints on mass, size, reliability, and mechanical complexity.

2.3 Sealed and Enclosed Wave Energy Conversion

A useful structural distinction in the literature is whether the power take-off (PTO) interacts directly with seawater or remains protected inside a rigid hull. Conventional wave energy converters include point absorbers, oscillating water columns, attenuators, overtopping devices, and other externally interacting systems. These concepts can be appropriate for large-scale power production, but exposed joints, hydraulic components, seals, turbines, and moving interfaces introduce maintenance and survivability concerns when the goal is a small, low-maintenance monitoring buoy [8], [9].

Rigid-hull and enclosed inertial WECs address this problem by placing the energy-conversion mechanism inside the buoy. Maheen and Yang review this class of devices and identify hull encapsulation as a way to protect the power take-off from direct marine exposure [8]. Qin et al. similarly motivate enclosed inertial WECs by noting that external WEC components are exposed to corrosion, biofouling, flooding, and extreme waves, whereas encapsulated devices harvest energy through relative motion between the buoy and an internal mass body [9]. This makes sealed internal conversion an important design direction for autonomous buoy applications.

The main trade-off is that sealing the mechanism removes direct access to wave force. The device must instead rely on the buoy's motion to excite an internal oscillator, pendulum, sliding mass, rolling mass, or gyroscopic reference. The resulting design problem is therefore to obtain sufficient internal relative motion while keeping the mechanism compact, robust, and compatible with electrical power generation.

2.4 Electromagnetic Harvesters for Buoys and Internal WECs

Electromagnetic conversion is one of the main candidate transduction methods for sealed buoy harvesters because it can be implemented using permanent magnets, coils, and mechanically simple relative motion [6], [10]. In linear electromagnetic devices, performance depends on the relative velocity between magnets and coils, the length over which the magnetic field interacts effectively with the winding, winding resistance, magnetic loading, and electrical-load matching. A useful benchmark within this group is the Teledyne work by Cheung and Childress, who developed a mass-spring-based low-frequency linear generator integrated with a spar buoy for ocean-monitoring applications. Their report combines sealed marine packaging, buoy integration, modelling, wave-tank testing, and field testing, with measured output from approximately 1.6 W to more than 6 W under the tested wave conditions [5]. It demonstrates that sealed linear electromagnetic buoy harvesting can move beyond the milliwatt scale.

Linear tubular permanent-magnet generators have also been studied for the Inertial Sea Wave Energy Converter (ISWEC). Bracco et al. designed and tested linear tubular generators for a gyroscopic inertial WEC concept [11], and Cappelli et al. further developed the corresponding machine design and experimental validation [12]. These studies are not direct templates for a small monitoring buoy because ISWEC uses a gyroscopic internal ref-

erence and targets a much larger power scale. Nevertheless, they are relevant because they show how internal relative motion in a sealed buoy can be coupled to tubular permanent-magnet generation.

Other compact electromagnetic buoy harvesters use different internal mechanisms. Li et al. proposed an electromagnetic ocean-wave energy harvester in which a swing body senses ultra-low-frequency wave motion and drives a high-speed electromagnetic power module through transmission gears [1]. This work is close in application because the device is installed inside a buoy and is intended for self-powered buoy operation. It represents a pendulum/rotary internal electromagnetic approach, where mechanical transmission is used to convert low-frequency buoy motion into higher-speed generator motion. Lo Schiavo et al. studied a rolling-magnet architecture, where a magnetic cylinder rolls along a track and induces voltage in coils connected to a DC load through a bridge rectifier [2]. This design is relevant as a low-cost data-buoy harvester and represents a mechanically simple rolling inertial approach.

The broader linear-generator literature reinforces both the promise and the limitations of electromagnetic power take-off. Ahamed et al. review linear-generator PTO systems for WECs and note that direct-drive linear machines can reduce mechanical complexity relative to hydraulic or turbine-type PTO systems, but also face challenges related to low-frequency operation, efficiency, magnetic loading, cogging, end effects, size, and cost [10]. These issues are directly relevant to compact electromagnetic harvesters. Direct electromagnetic conversion can reduce transmission complexity, but useful power still depends on sufficient relative velocity, active interaction length, magnetic-field strength, coil resistance, and suitable electrical loading.

2.5 Alternative Transduction Technologies

Triboelectric and hybrid nanogenerator systems form another group of small-scale wave-energy harvesters. Triboelectric systems generate charge through repeated contact, separation, or sliding between materials, while hybrid nanogenerators combine more than one transduction mechanism, for example triboelectric and electromagnetic conversion, to make use of different parts of the motion response [3], [4]. These approaches are relevant because low-frequency and irregular wave motion can be converted into repeated contact, pendulum, rolling, or rotational motion through an internal motion-conditioning mechanism rather than through direct coupling to a conventional rotary generator.

Jung et al. developed a self-powered ocean buoy using a disk-type triboelectric nanogenerator with a mechanical frequency regulator [3]. In that system, slow wave-induced motion is converted into stored rotational motion and then released at higher speed through the frequency-regulation mechanism. The design addresses a common problem in wave-driven nanogenerators: the excitation is slow and irregular, while the transducer benefits from faster repeated mechanical interaction. Zhao et al. presented a pendulum-based hybrid nanogenerator for high-entropy wave energy harvesting, combining triboelectric and electromagnetic generation with chaotic pendulum motion and power-management circuitry [4]. This type of design addresses broadband and multi-directional excitation by allowing the internal pendulum response to become more complex than a single tuned linear oscillation.

These studies show that the transduction mechanism and the motion-conditioning mechanism are often closely linked. They also introduce practical design issues that differ from those of permanent-magnet electromagnetic generators, including high-voltage/low-

current output, impedance matching, rectification and storage, material wear, contact durability, and environmental sealing. For the present thesis, they are relevant because they show alternative ways of harvesting low-frequency irregular buoy motion and help define the broader design space considered before the architecture selection in Chapter 3.

2.6 Low-Frequency and Irregular-Wave Motion Conditioning

The most persistent technical challenge across the literature is the low-frequency and irregular character of wave excitation. A simple linear resonant harvester can perform well near its natural frequency, but ocean waves are broadband and time-varying. If the device is too narrowly tuned, much of the available motion is not converted effectively.

Several strategies have been proposed to address this mismatch. Variable-stiffness and negative-stiffness mechanisms can shift or broaden the response of a wave energy converter. Dang et al. proposed a direct-drive WEC using a variable stiffness mechanism and pre-tension springs to improve the motion and output of a supplementary mass coupled to a linear generator [13]. Qin et al. used a magnetic nonlinear stiffness mechanism in an enclosed inertial WEC and reported improved response and wider frequency bandwidth under low-frequency waves compared with a linear-stiffness version [9]. These studies show that the restoring-force system can be a major design lever for shifting internal harvester motion toward the wave-frequency range.

Other approaches mechanically transform the wave-induced motion before it reaches the generator. Ummaneni et al. proposed using gas springs to convert low-speed, high-force wave motion into higher-speed, lower-force motion more suitable for a linear machine [14]. Chen et al. introduced a wave-roller mechanical interface intended to reduce the required mover stroke and allow multiple linear generators to operate from the same captured wave motion [15]. These concepts are not architecture selections for the present work, but they highlight an important design principle: for compact generators, mechanical conditioning of the excitation may be required before efficient electrical conversion is possible.

Nonlinear and multistable pendulum systems offer another route to broader response. Wang and Wang proposed a gimbaled-pendulum harvester with magnetic multistability for multi-directional irregular excitation and reported increased average output under certain motion directions, together with field results from a compact prototype [16]. Crowley et al. modelled the WITT wave energy converter, an inertial pendulum-based concept whose response can involve multiple modes rather than a single simple resonance [17]. Zhao et al. similarly present pendulum dynamics as a way to convert broadband, disordered wave excitation into useful generator motion [4]. These papers show why sealed internal harvesters should be evaluated under irregular and broadband excitation rather than only under ideal sinusoidal input.

2.7 Generator Design, Sealing, and Electromagnetic Detail

Several generator-focused papers are relevant mainly as supporting design references rather than as direct architecture matches. Jang et al. analysed and experimentally verified a

moving-magnet linear actuator with a cylindrical Halbach array, illustrating how cylindrical magnet arrangements can be treated analytically and experimentally in linear electromagnetic machines [18]. Farrok et al. proposed a permanent-magnet linear generator for ocean-wave conversion with attention to demagnetisation resistance, an issue that becomes more important as magnetic loading and power density increase [19]. Crozier et al. presented a direct-drive wave energy converter system based on a linear generator, providing a useful example of coupled electromechanical design and testing for wave energy conversion [20]. Liu et al. proposed an oscillating-water-column WEC with a linear generator, magnetic-fluid reciprocating seal, and air spring, which is not a sealed internal inertial buoy harvester but is useful as an example of how sealing and linear PTO design remain central problems in marine generators [21].

These generator-focused studies show that marine linear generators are constrained by both electromagnetic and mechanical requirements. Higher magnetic loading, longer active regions, and smaller air gaps can improve conversion, but they also affect size, moving mass, stroke, sealing, and structural layout. For compact sealed buoy harvesters, the generator geometry therefore has to be considered together with the internal motion mechanism and the available package volume.

2.8 Modelling and Experimental Validation Approaches

Wave energy systems are commonly modelled at different levels of fidelity, from simplified single-degree-of-freedom models to coupled hydrodynamic and power-take-off simulations. For a sealed internal buoy harvester, the most relevant modelling task is to connect buoy or base motion to internal relative motion, generator voltage, electrical damping, and load power. Full buoy hydrodynamics remains important for a deployable system, but reduced-order models are useful when the immediate objective is to test an internal generator concept and identify its scaling tendencies.

Katsidoniotaki et al. discuss mid-fidelity model verification for a point-absorbing WEC with linear PTO, showing the need to check reduced-order simulations against higher-fidelity or experimental references [22]. The surface-riding WEC study by Jin et al. provides another example in which hydromechanical response and linear-generator output are evaluated together [23]. These studies support the use of experimental comparison before using a reduced-order model for design interpretation or scaling.

2.9 Local Icelandic and Reykjavík University Context

The present work continues a local research line at Reykjavík University on self-powered buoys and wave-driven power generation. Guðmundsson studied electrical power generation for a whale-monitoring buoy and analysed Icelandic wave conditions as part of estimating the production capability of a buoy-mounted generator [24]. That work connects the technical problem to Icelandic operating conditions and to the practical need for a buoy that can power onboard electronics.

Olgeirsson later investigated a self-sustaining power supply for a low-power sea buoy, comparing possible ocean-energy sources and developing a prototype-level electromagnetic power concept for low-power operation [25]. Together, these projects establish the local application context: a buoy-scale power supply, driven by Icelandic wave conditions,

intended to reduce dependence on finite onboard energy storage. The remaining step addressed in this thesis is a more detailed evaluation of sealed internal harvester architectures, including electromagnetic modelling, physical prototyping, experimental characterisation, and model-based scaling.

2.10 Research Gap and Thesis Scope

The reviewed literature shows three points that are central to this thesis. First, sealed or enclosed internal WECs are attractive when survivability and low maintenance are more important than direct access to wave force. Second, internal harvesters must deal with low-frequency, irregular excitation, which often requires compliance, frequency conversion, nonlinear motion, or other motion-conditioning strategies. Third, there is a gap between low-power buoy harvesters aimed mainly at sensor-scale loads and larger enclosed WEC concepts developed around more specialised or higher-power systems [2], [3], [4], [5].

The design space of interest is therefore a sealed internal harvester that is simple enough to build and test, but still scalable through physical quantities that can be modelled. Chapter 3 uses this design space as the basis for comparing candidate internal architecture families and selecting the architecture developed in the remainder of the thesis.

Chapter 3

Methodology and Design Concept

3.1 Design Process Overview

The project follows a standard engineering design process combining numerical modelling with physical testing. The work is divided into four main phases:

1. Design and Concept Selection.
2. Numerical Modelling.
3. Prototyping and Testing.
4. Scaling and Optimisation.

3.2 Concept Selection

This section addresses Research Question 1: *Which internal mechanical architecture is most suitable for this thesis given the requirements of fully sealed operation, irregular-wave compatibility, practical implementation, and future scaling potential?* A structured concept-selection process was used to compare candidate architecture families against the main project requirements, so that the reasoning behind the final choice is explicit and traceable rather than dependent on qualitative discussion alone.

3.2.1 Scope of Considered Concepts

Sealed internal operation is treated here as a hard requirement rather than as one of several criteria to be weighed against the others. Architectures that depend on exposed moving interfaces, hull-crossing transmission elements, or power take-off mechanisms in direct contact with seawater are excluded from the comparison at the outset, regardless of their other merits. This filter rules out conventional point absorbers, hydraulic-piston devices, and other externally interacting wave-energy converters, and restricts the comparison to architectures whose main energy-conversion mechanism remains fully enclosed [8].

3.2.2 Architecture Families Considered

Within the sealed-internal scope, four candidate architecture families are considered, chosen to span the principal mechanical types reported in the literature for harvesting wave-induced motion inside a buoy:

1. **Linear sliding architecture:** a translating internal mass, magnetic assembly, or coil assembly moving along a guided linear path relative to the buoy structure and coupled directly to a linear power take-off.
2. **Pendulum or swing-body architecture:** an internal oscillating mass that converts buoy motion into rotational or reciprocating generator motion, either directly or through internal transmission elements.
3. **Rolling inertial architecture:** an internal rolling magnetic or mechanical mass moving along a curved or guided track inside the buoy.
4. **Gyroscopic inertial-reference architecture:** an internal rotating reaction body that creates relative motion or reaction torque within the buoy and drives a power take-off system.

3.2.3 Cross-Cutting Dynamic Strategies in the Literature

Nonlinear and variable-stiffness mechanisms are not treated as a separate top-level architecture family. They are cross-cutting dynamic strategies that can be incorporated into multiple families to broaden useful response under irregular wave excitation. Pendulum-based systems may be made multistable through magnetic restoring forces, and enclosed inertial converters may use nonlinear stiffness to increase bandwidth or improve energy capture under stochastic motion [9], [16]. These approaches address the low-frequency and irregular nature of ocean waves, but they do not by themselves define a unique mechanical family independent of the underlying architecture.

3.2.4 Representative Reference Concepts from the Literature

The candidate families are grounded in representative concepts reported in the literature. Linear sliding and direct-drive internal electromagnetic concepts are represented by the sealed low-frequency generator systems developed for ocean sensing and buoy applications [5]. Pendulum and swing-body concepts are represented by internal oscillating-mass harvesters for buoy-scale or wave-driven energy conversion [1], [17]. Rolling inertial concepts are represented by rolling magnetic-cylinder harvesters for small data buoys [2]. Gyroscopic inertial-reference systems are represented by ISWEC-type devices, in which an internal rotating body provides the inertial reference for energy conversion [12].

These references define the comparison space, but the comparison is made at the architecture family level rather than as a detailed optimisation of each published device. Linear sliding systems offer the most direct electromagnetic scaling route because active length, coil count, magnetic coupling, moving mass, and electrical loading can be varied within a one-dimensional model. Pendulum and swing-body systems are well suited to low-frequency excitation because their effective natural frequency can be reduced through geometry and inertia, but they often require rotational transmission, gearing, or multi-degree-of-freedom modelling when coupled to an electromagnetic generator. Rolling inertial systems are mechanically simple and attractive for sealed low-cost data buoys, but the reported systems are mainly aimed at low-power sensor-scale applications and provide less direct scaling through active electromagnetic interaction length. Gyroscopic inertial-reference systems can reach higher power levels in larger enclosed WECs, but their rotating machinery, control requirements, and coupled dynamics make them less suitable for a low-cost bench-scale thesis prototype.

The goal is therefore not to reproduce any single published design, but to identify the architecture family most compatible with the thesis objectives, available components, modelling scope, and prototype constraints.

3.2.5 Evaluation Criteria and Decision Matrix

Sealed internal operation is treated as a precondition (Section 3.2.1) and therefore does not appear in the matrix. The remaining criteria reflect the priorities of a self-powered monitoring buoy that must demonstrate a credible path toward the 10 W target, be analysable within the timeframe of the thesis, and be inexpensive enough to prototype and iterate. Each architecture is scored against the criteria in Table 3.1 on a five-point scale, where a higher score indicates better suitability.

Table 3.1: Evaluation criteria used to compare candidate internal architecture families.

Criterion	Weight	Rationale
Power output and scaling potential	0.25	The architecture should offer a credible path from prototype scale toward the 10 W target through increased active length, improved interaction, or modular expansion.
Modelling and analytical tractability	0.20	The architecture should be amenable to analytical and numerical modelling within the timeframe of the thesis, so that a predictive model can be developed and used to support scaling analysis.
Mechanical simplicity and reliability	0.20	Lower complexity reduces failure points and improves long-term robustness in a sealed deployment.
Manufacturability, cost, and prototyping practicality	0.20	The architecture should be practical and inexpensive to fabricate, modify, and re-test within the budget, facilities, and components available for the thesis.
Suitability for low-frequency irregular motion	0.15	The architecture should respond credibly to broadband stochastic buoy motion. This is weighted lower than the other criteria because nonlinear or compliance-based modifications to a resonant architecture are an explicit subject of investigation in this thesis.

The weighted total score for each architecture is calculated as

$$S = \sum_{i=1}^n w_i x_i \quad (3.1)$$

where w_i is the weighting factor for criterion i and x_i is the corresponding score assigned to the architecture under evaluation. The resulting matrix is shown in Table 3.2.

Table 3.2: Weighted decision matrix for the candidate internal architecture families. Scores are on a 1–5 scale; criteria C1–C5 correspond to those in Table 3.1 in the same order.

Architecture	C1	C2	C3	C4	C5	Total
Linear sliding	5	5	4	5	3	4.50
Rolling inertial	2	4	5	5	3	3.75
Pendulum / swing-body	3	3	3	3	4	3.15
Gyroscopic inertial-reference	4	2	2	2	4	2.80

The scoring reflects the following observations from the reviewed literature and from the practical constraints of the project. On power potential and scaling, sealed linear electromagnetic generators have demonstrated continuous outputs in the range of several watts at bench and field scale [5], providing a more direct path toward the 10 W target than the milliwatt-scale rolling and pendulum-based systems reported for small drifters and IoT sensors [2]. Gyroscopic systems can in principle reach much higher power, but only at scales well beyond the man-portable buoy considered here [12]. On modelling tractability, the linear architecture reduces naturally to a one-dimensional electromagnetic and mechanical problem, while gyroscopic and geared pendulum systems require multi-degree-of-freedom dynamic models that are difficult to develop and validate within the thesis timeframe. On simplicity, manufacturability, and cost, rolling-inertial and linear designs are favoured by the use of off-the-shelf magnets, simple guides, and standard coil-winding techniques. On low-frequency suitability, pendulum and gyroscopic concepts have a structural advantage because long pendulums and large rotating bodies naturally produce low effective frequencies, while linear and rolling systems must rely on compliance design or nonlinear restoring elements to operate at the low frequencies typical of ocean waves.

3.2.6 Selection Outcome

The linear sliding architecture obtains the highest weighted score and is selected as the baseline for the remainder of the project. The decision is driven by three factors: a credible scaling path toward higher average power, supported by published sealed linear generators operating at the watt scale; a relatively direct route to analytical and numerical modelling within the thesis timeframe; and low mechanical and component cost compared to gyroscopic or geared pendulum alternatives. Suitability for low-frequency irregular excitation is the principal weakness of this family and is addressed directly through system-level model, the prototype experiments, and the coupled validation developed in Chapters 6–8.

The rolling inertial architecture scores second and remains a credible alternative if future iterations prioritise sealing simplicity and structural robustness over electromagnetic interaction length. Pendulum and gyroscopic concepts are judged less tractable to model and prototype at the scale and timeframe of this project.

A further practical factor supporting the selected architecture was the availability of an existing segmented linear-motor magnetic rod at Reykjavík University. This component did not determine the architecture by itself; the primary reasons for the selection remain the sealed internal configuration, mechanical simplicity, modelling tractability, and scaling potential identified in the decision matrix. However, the rod provided a high-quality periodic magnetic assembly with sufficient active length for a bench-scale prototype. Its availability reduced cost and development risk, and allowed the thesis work to focus on

electromagnetic characterisation, coil design, system modelling, experimental testing, and coupled validation rather than on the manufacture of a custom magnet array.

3.3 Selected Architecture

The selected linear sliding architecture generates electrical energy through linear relative motion between a periodic magnetic assembly and a guided coil assembly. In the present thesis it is implemented using the available segmented magnetic rod and a translating coil carriage constrained to move along the rod axis. The rod provides a spatially varying magnetic field, and the carriage moves the active coils through this field to induce voltage. The operating principle is therefore internal relative linear motion within a sealed enclosure, rather than wave-driven motion transmitted through an external power take-off.

The selected architecture follows the same general principle as other enclosed linear electromagnetic buoy harvesters, where useful power is produced from relative motion between an internal inertial subsystem and a magnetic power take-off. The main design knowledge taken from these systems is that sealed linear electromagnetic generation can produce useful watt-scale power when sufficient stroke, relative velocity, active magnetic length, and electrical loading are available. At the same time, the literature highlights the main weaknesses of this architecture: low-frequency wave excitation can produce limited relative velocity, the required stroke can become large, and the generator geometry must be scaled without making the device too long or mechanically fragile. These issues are therefore carried forward explicitly in this thesis through the restoring-system design, the rod-and-coil electromagnetic model, the prototype tests, and the later stroke and scaling analysis.

The architecture is modular along its main axis. Active length can be increased by extending the translating assembly or by adding further coils, which provides a clear scaling lever for later analysis, and the same basic principle can be studied at bench scale before being connected to more realistic buoy-motion conditions. Detailed geometric characterisation and electromagnetic modelling of the selected rod-and-coil architecture are presented in Chapter 4, while the physical prototype and measurement system are described in Chapter 5.

Chapter 4

Electromagnetic Modelling and Coil Optimisation

4.1 Modelling Scope and Strategy

After the linear sliding architecture was selected in Section 3.3, the next step was to develop a quantitative electromagnetic model of the segmented magnetic rod and the translating coil assembly. The prototype used an existing segmented magnetic rod as the field source, which avoided the need to manufacture a custom magnet array. Since the rod was not supplied with a detailed datasheet, the magnet and spacer geometry had to be determined from measurement, while the effective magnetic strength was calibrated from the measured surface field.

This chapter focuses on identifying the rod geometry, reconstructing the magnetic field produced by the rod, and determining how that field couples to a translating coil. The modelling workflow consisted of three main stages:

1. **Rod characterisation:** the segmented magnetic rod was identified, its external dimensions were measured, and its surface magnetic field was scanned with a Hall probe. The measured field pattern was then used to infer the internal magnet–spacer geometry required for modelling.
2. **Finite-element field reconstruction:** the rod field was reconstructed using two independent axisymmetric magnetostatic FEM models. ANSYS Maxwell was used first as a commercial-solver reference, while MATLAB PDE Toolbox was then used to model the full rod and provide the field maps used in the later MATLAB-based analysis.
3. **Coil-level optimisation:** the validated field maps were used to compute flux linkage, induced voltage, and matched-load power for candidate coil geometries. The axial winding length was fixed from the magnetic rod pitch to avoid cancellation between opposing field regions, while the coil outer diameter was selected from the optimisation sweep.

The coil model used in this chapter treats the prototype coils as air-core windings. Ferromagnetic cores, laminated flux concentrators, or other soft-magnetic structures around the coils were not included in the optimisation. Such components could increase the effective flux linkage and electromagnetic coupling, but they would require a separate magnetic-circuit design and would introduce additional issues such as saturation, hysteresis, eddy-

current loss, cogging force, and manufacturing complexity. Custom laminated coil cores were therefore left outside the prototype scope and are treated as a possible route for future electromagnetic redesign.

The result is a measured-field-consistent model of the segmented magnetic rod and a coil-geometry recommendation for the physical prototype described in Chapter 5. This chapter is limited to the electromagnetic part of the problem: how the magnetic field links with the coil, what voltage is induced by relative motion, and how the coil geometry affects the predicted matched-load power. The full electromechanical response is developed in Chapter 6. The corresponding prototype experiments are presented in Chapter 7, the coupled model is validated against those measurements in Chapter 8, and the validated model is then used for the scaling analysis in Chapter 9.

4.2 Magnetic Rod Characterisation as Model Input

The electromagnetic model required the rod dimensions, axial segmentation, and effective magnetic strength. These inputs were obtained from the rod markings, external measurements, surface-field scans, and interpretation of the measured field.

4.2.1 Rod Identification and External Geometry

The rod was identified from its markings as a LinMot linear-motor slider originally manufactured for Valka, with part number 0150–2889 and family designation PL01-24x1010/930. The measured total length was approximately 1010 mm, and the measured outer diameter was approximately 27 mm. The family designation “PL01-24” suggests a nominal internal magnetic diameter of approximately 24 mm, but the detailed magnet and spacer arrangement was not available.

These measurements defined the overall rod envelope used in the electromagnetic models. The remaining inputs were the internal magnet–spacer geometry and the effective magnetic strength of the assembly. These were inferred and calibrated from the surface-field measurements described in the next section.

4.2.2 Surface Magnetic Field Measurements

The rod surface field was measured using a DC gaussmeter with a Hall-effect probe. Measurements were taken manually over one observed magnetic repeat interval of approximately 60 mm, with the probe placed against the rod surface. Both radial and axial field components were recorded.

Two independent passes were carried out to check repeatability. Figure 4.1 compares the raw scans. Their close agreement supports using the measured field as the calibration target for the analytical and FEM models.




Figure 4.1: Comparison of two raw surface-field measurement passes over one magnetic repeat period.

The measured profile is consistent with a bucking segmented magnet assembly. The radial component forms paired local maxima around each inferred spacer region, with the largest peaks reaching approximately 7100 G. The radial field at the midpoint between each peak pair was approximately 6200 G. Between neighbouring spacer regions, the radial component decreased towards the centre of each magnet body, while the axial component remained approximately 3000 G to 3500 G.

Measurements near the physical rod ends showed lower radial peaks of approximately 3200 G to 4000 G, consistent with weaker flux concentration where the end magnets do not have opposing-polarity neighbours on both sides. The measured field profiles were therefore used to infer the axial segmentation and calibrate the magnetic field models developed in the following sections.

4.2.3 Field-Derived Rod Geometry

The internal axial geometry was inferred from the radial surface-field profile. In a bucking assembly, adjacent magnets are magnetised in opposite axial directions, and the strongest radial field occurs near the magnet-spacer transitions, where flux is driven outward between opposing-polarity magnet sections. The peak locations were therefore used to estimate the magnet and spacer lengths.

The short spacing between the two peaks surrounding one spacer region was approximately 7 mm and was interpreted as the spacer length. The longer spacing between neighbouring spacer regions was approximately 23 mm and was interpreted as the magnet length. This gives a structural pitch of

$$p_{\text{struct}} = 23 \text{ mm} + 7 \text{ mm} = 30 \text{ mm}. \quad (4.1)$$

Because the magnetisation direction alternates between neighbouring magnets, the magnetic polarity pattern repeats after two structural sections. The magnetic repeat period is therefore

$$p_{\text{mag}} = 2p_{\text{struct}} = 60 \text{ mm.} \quad (4.2)$$



Figure 4.2: CAD visualisation of the inferred segmented magnetic-rod geometry. Red and blue indicate opposite magnetisation directions; grey indicates spacers and end regions.

The rod contained 31 magnets in total, as established from direct measurement. The geometric quantities used in the model are summarised in Table 4.1.

Table 4.1: Segmented magnetic-rod geometry used in the model.

Component	Dimension	Value	Basis
Rod assembly	Total length	1010 mm	Measured
	Outer diameter	27 mm	Measured
	Total magnet count	31	Measured
Magnet	Axial length	23 mm	Field-derived
Spacer	Axial length	7 mm	Field-derived
Axial pattern	Structural pitch	30 mm	Derived
	Magnetic repeat period	60 mm	Derived from bucking polarity pattern
End regions	Non-magnetic end 1	33 mm	Measured
	Non-magnetic end 2	54 mm	Measured

This geometry was used as the axial structure for the analytical and finite-element models.

4.2.4 Analytical Interpretation of the Measurements

The measured field was first interpreted using a MATLAB analytical model based on the Biot–Savart field of circular current loops. A uniformly magnetised cylindrical magnet can be represented by an equivalent bound surface current, so each magnet was approximated as a finite stack of circular current loops with equivalent surface current density

$$K = M = \frac{B_{\text{rem}}}{\mu_0}, \quad (4.3)$$

where B_{rem} is the remanence parameter used in the analytical model. Contributions from neighbouring magnets were added with alternating polarity to reproduce the bucking arrangement, and a three-point axial average over 0.5 mm was applied to approximate the finite probe width.

The analytical model was used to test how the inferred axial pattern depended on the radial position of the magnetic material. The PL01-24 designation suggests a 24 mm magnetic stack inside the measured 27 mm outer diameter, corresponding to a tube wall thickness of approximately 1.5 mm. This thick-wall interpretation produced a radial-field profile that was too smooth near the magnet–spacer transitions. Increasing the non-magnetic wall

thickness increases the distance between the magnet surface and the probe, which attenuates sharp spatial features in the measured field. A larger magnetic radius of approximately 13.0 mm, corresponding to a wall thickness of approximately 0.5 mm, reproduced the measured transition-region gradients more closely.




Figure 4.3: Analytical comparison of two radial geometry interpretations. The thick-wall interpretation smooths the transition peaks too strongly, while the thin-wall interpretation better reproduces the measured radial-field gradients near the magnet–spacer transitions.

The thin-wall interpretation was used for the explicit ANSYS rod geometry. In the MATLAB PDE model used later in this chapter, the rod was instead represented by an effective magnetic field model tuned directly to the measured surface field. For the coil-optimisation step, the important quantity is the external field in the winding region, which is constrained by the measured field calibration.

4.3 Finite-Element Field Reconstruction

The magnetic rod was then modelled using two finite-element magnetostatic approaches. The ANSYS Maxwell model provided an independent commercial-solver reference, while the MATLAB PDE Toolbox model provided the full-rod field reconstruction used as the primary model for later analysis.

4.3.1 ANSYS Maxwell Model

The first FEM model was built in ANSYS Maxwell 2D using an axisymmetric r – z magnetostatic formulation. The model was generated through Python automation using PyAEDT. The script created the rod geometry, assigned magnetisation directions and material properties, defined the surrounding solution domain, solved the model, and exported the magnetic field on a regular r – z grid for later MATLAB processing.

The ANSYS model used the thin-wall geometry suggested by the analytical interpretation: a magnetic core radius of 13.0 mm, a tube wall thickness of 0.5 mm, and a resulting outer radius of 13.5 mm. The magnet relative permeability was set to $\mu_r = 1.05$, and the magnet coercivity was tuned so that the exported surface field matched the measured radial-field peak.

A practical limitation was the mesh-size cap in the ANSYS Electronics Desktop Student edition. The full 1010 mm, 31-magnet rod could not be meshed at sufficient resolution while also resolving the sharp gradients near the magnet–spacer transitions. The ANSYS model was therefore limited to a representative 7-magnet subsection. This subsection was

sufficient to capture the local periodic field structure relevant to a single coil, but it could not represent full-rod end effects.




Figure 4.4: Magnetic field-line plot from the ANSYS Maxwell model of a representative 7-magnet subsection. The alternating bucking magnetisation concentrates flux near the magnet-spacer transition regions.

The solved field was exported over the winding region used in the coil-level calculations. The export covered radial positions from the rod surface outward into the candidate coil region and an axial window centred on a representative spacer.

4.3.2 MATLAB PDE Toolbox Model

The second FEM model was built in MATLAB PDE Toolbox using the `femodel` magneto-static interface in 2D axisymmetric mode. This removed the ANSYS Student-edition mesh limitation and allowed the full 1010 mm rod, including all 31 magnets and both end regions, to be solved in one model.

In a current-free axisymmetric region with permanent magnets, the governing magnetostatic equation may be written in terms of the magnetic vector potential as

$$\nabla \times \left(\frac{1}{\mu_0 \mu_r} \nabla \times \mathbf{A} \right) = \nabla \times \mathbf{M}, \quad (4.4)$$

where $\mathbf{M}$ is the prescribed magnetisation and μ_r is the relative permeability of each material region. In the axisymmetric model, each magnet was assigned a uniform axial magnetisation with alternating sign between neighbouring magnets.

The MATLAB PDE model used for the later calculations was an effective-field model rather than an explicit tube-wall model. The magnetic region was extended to the measured rod radius, and the remanence input was tuned to match the measured surface field. The exported effective-field case used $B_{\text{rem,eff}} = 1.46$ T and $\mu_r = 1.05$. A separate thin-wall case with a 13.0 mm magnetic core and 0.5 mm tube wall was kept as a reference case, while the results presented in the remainder of this chapter use the tuned effective-field export.

Two MATLAB PDE field maps were exported. The first was a fine-resolution coil-window map centred on a representative spacer and used for field validation and coil OD optimisation. The second was a coarser full-rod map used to visualise the global field structure.




Figure 4.5: Radial magnetic field component $B_r(r, z)$ over the full 1010 mm rod from the MATLAB PDE Toolbox model. The periodic bucking structure is visible along the rod, with reduced field amplitude near the end regions.




Figure 4.6: Radial and axial magnetic field components along a line just outside the rod surface, evaluated over the full rod length. The central region is nearly periodic, while the end regions show the expected reduction in field amplitude.

The optimum coil OD is determined from the representative central field map, not from the full-rod visualisation. The full-rod results are included to show the periodic field structure over the complete rod and the reduced field amplitude near the end regions.

4.4 Field-Map Calibration and Agreement

Both FEM field maps were compared with the measured Hall-probe data over one 60 mm magnetic repeat period. The purpose of this comparison was to check that the calibrated models reproduce the measured field structure at the surface-probe radius before using the field maps for coil calculations.

Figure 4.7 compares the measured radial and axial field components with the ANSYS and MATLAB PDE field maps.




Figure 4.7: Comparison between measured and modelled field components over one magnetic repeat period at the surface-probe radius. The radial component shows the paired peak structure used to infer the magnet–spacer geometry, while the axial component provides an additional check of the field variation through the magnet bodies.

The two FEM models use different internal representations of the rod, but both are calibrated to the same measured surface field. The ANSYS model uses an explicit thin-wall geometry and is tuned through the magnet coercivity. The MATLAB PDE model uses an effective magnetic geometry and is tuned through the remanence input. These calibrated inputs should not be interpreted as independent measurements of the physical magnet grade; they are model parameters selected to reproduce the measured external field.

The numerical agreement is summarised in Table 4.2. The radial-field RMSE was approximately 0.071 T for ANSYS and 0.061 T for MATLAB PDE. The axial-field RMSE was approximately 0.065 T for both models. The remaining differences were considered acceptable for the coil-design calculations, which depend mainly on the spatial field structure and its radial decay through the winding region.

Table 4.2: Field-level calibration summary against measured Hall-probe data over one magnetic repeat period.

Field map	Radial RMSE	Axial RMSE
ANSYS Maxwell, 8-magnet subsection	0.071 T	0.065 T
MATLAB PDE Toolbox, full-rod model	0.061 T	0.064 T

Both field maps reproduce the main measured features: the periodic bucking structure, the radial peaks near the magnet–spacer transitions, and the axial-field plateaus inside the magnet bodies. They were therefore used for the coil-geometry comparison in the following sections.

4.5 Coil-Level Electromagnetic Calculations

The calibrated field maps were processed in MATLAB to estimate flux linkage, induced voltage, and matched-load power for candidate coil geometries. The calculations in this section are performed for a single coil, and the resulting geometry recommendation is applied to each coil in the prototype.

4.5.1 Flux Linkage from the Vector Potential

In an axisymmetric geometry, the magnetic flux through a circular turn of radius r at axial position z can be written in terms of the azimuthal vector potential $A_\phi(r, z)$ as

$$\Phi(r, z) = 2\pi r A_\phi(r, z). \quad (4.5)$$

The vector potential was reconstructed from the exported radial field component using

$$B_r = -\frac{\partial A_\phi}{\partial z}. \quad (4.6)$$

Numerically, this gives

$$A_\phi(r, z) = -\int_{z_{\min}}^z B_r(r, \zeta) d\zeta + C(r), \quad (4.7)$$

where the integration constant $C(r)$ does not affect the induced voltage because only the axial derivative of the total flux linkage is used.

For a multi-turn coil, the total flux linkage at carriage offset z_0 was obtained by summing the contribution from all turn-centre locations:

$$\lambda(z_0) = \sum_{i=1}^{N_{\text{turns}}} 2\pi r_i A_\phi(r_i, z_i + z_0). \quad (4.8)$$

4.5.2 Induced Voltage and Matched-Load Power

For relative motion between the coil and the rod at constant speed v , Faraday's law gives

$$e(z_0) = -\frac{d\lambda}{dt} = -v \frac{d\lambda(z_0)}{dz_0}. \quad (4.9)$$

The RMS induced voltage was computed over one 60 mm magnetic period. For design comparison, the maximum average power into a matched resistive load was approximated as

$$P_{\max} = \frac{E_{\text{rms}}^2}{4R_{\text{coil}}}, \quad (4.10)$$

where R_{coil} is the estimated DC resistance of the coil winding.

These calculations are used only for comparing coil geometries. They assume constant relative speed and matched resistive loading, so they do not include the coupling between electrical damping and carriage motion. That coupled response is treated later in the system-level model.

4.6 Coil Outer-Diameter Optimisation

The coil geometry was reduced to two main design choices: the axial winding length and the radial winding thickness.

The axial winding length was set by the magnetic pitch of the rod. As shown in Section 4.2.3, the rod has a structural pitch of 30 mm and a magnetic repeat period of 60 mm. If the coil is made much longer than one structural pitch, it begins to span neighbouring

regions of opposite magnetic phase. Additional turns then contribute partly opposing voltage rather than increasing the output proportionally.

The active winding length was therefore chosen to fit within one 30 mm structural pitch. Each prototype coil was wound on a 3D-printed reel with a total length of 30 mm and 1.5 mm end flanges, giving an active winding length of

$$L_{\text{coil}} = 30 \text{ mm} - 2 \times 1.5 \text{ mm} = 27 \text{ mm}. \quad (4.11)$$

The coil inner diameter was fixed by the clearance required for smooth carriage motion along the rod. Iterative testing showed that a reel inner diameter of 27.6 mm allowed reliable translation over the 27.0 mm rod without excessive friction. With a 0.8 mm central tube wall, this gave a fixed coil inner diameter of 29.2 mm.

With the axial winding length and inner diameter fixed, the remaining design variable was the coil outer diameter (OD). Increasing the OD allows more radial winding layers and therefore more turns, but it also increases the copper length and winding resistance. The optimum occurs where the additional flux-linkage benefit from extra turns is balanced by the resistance penalty.

A parametric sweep of candidate outer diameters was then conducted. The sweep evaluated buildable coil geometries using the wire and winding constraints of the prototype. The active axial winding length was fixed at 27.0 mm, the coil inner diameter at 29.2 mm, the packed wire diameter at 0.53 mm, and the copper diameter at 0.475 mm. The copper resistivity was taken as $\rho = 1.724 \cdot 10^{-8} \Omega \text{ m}$. All candidates were evaluated at a fixed relative speed of 0.5 m/s.

For the dual-FEM comparison, a fill factor of $f = 0.85$ was used. The fill factor enters through the effective wire-pack spacing $d_{\text{eff}} = d_{\text{pack}}/\sqrt{f}$, which determines how many radial layers and turns per layer fit within the winding window. Each candidate OD therefore corresponds to an integer number of layers and turns.

4.6.1 Cross-Validation Result

Figure 4.8 shows the predicted matched-load power as a function of coil outer diameter for both FEM field maps.




Figure 4.8: Predicted matched-load power P_{max} as a function of coil outer diameter, comparing the ANSYS Maxwell and MATLAB PDE field maps at $v = 0.5 \text{ m/s}$ and $f = 0.85$. Both models identify the same optimum at $\text{OD} \approx 48.7 \text{ mm}$.

Both field maps identify the same optimum coil outer diameter:

$$\text{OD}^* \approx 48.7 \text{ mm.} \quad (4.12)$$

The numerical results at this optimum are summarised in Table 4.3. The ANSYS and MATLAB PDE predictions differ slightly in absolute power, but both give the same optimum diameter.

Table 4.3: Coil OD optimisation summary for the two FEM field maps, evaluated at $v = 0.5 \text{ m/s}$ and $f = 0.85$.

Field map	OD*	$P_{\max}$	E_{rms}	R_{coil}
ANSYS Maxwell	48.75 mm	1.338 W	7.06 V	9.31Ω
MATLAB PDE Toolbox	48.75 mm	1.293 W	6.94 V	9.31Ω

The power values in Table 4.3 are matched-load predictions for a fixed relative speed and should not be interpreted as the expected wave-driven prototype output. For the coil design, the main result is the optimum diameter. The agreement between the two field maps indicates that this diameter is not strongly affected by the solver implementation or by the different internal rod representations used in the two models.

4.6.2 Sensitivity to Winding Fill Factor

The predicted power from the OD sweep depends on the assumed winding fill factor, since the fill factor affects both the number of active turns and the winding resistance. A sensitivity study was therefore performed to check whether the recommended coil diameter depended strongly on the assumed packing density.

The MATLAB PDE field map was re-analysed using fill factors of 0.55, 0.65, 0.75, and 0.85. The results are shown in Figure 4.9.




Figure 4.9: Sensitivity of predicted matched-load power to winding fill factor using the MATLAB PDE field map. The predicted power level increases with fill factor, while the optimum coil outer diameter shifts by less than 0.5 mm across the tested range.

As expected, higher fill factors increase the predicted matched-load power by allowing more copper to fit within the same winding volume. However, the optimum OD shifts

by less than 0.5 mm across the tested range. This shift is smaller than the packed wire diameter, so the recommended coil diameter is not sensitive to the exact packing density achieved during winding.

4.6.3 Recommended Coil Outer Diameter

Combining the dual-FEM comparison in Section 4.6.1 with the fill-factor sensitivity study in Section 4.6.2, the recommended coil outer diameter for the modelled rod and winding constraints is

$$\boxed{\text{OD}^* \approx 48.7 \text{ mm}}. \quad (4.13)$$

This is the electromagnetic design recommendation for the coil geometry considered in this chapter. It is based on field-level modelling at fixed relative speed with matched resistive loading, and was selected to maximise the predicted coil-level power. At this OD and a fill factor of $f = 0.85$, the optimisation corresponds to approximately 782 turns. For the prototype, the target was rounded to 800 turns per coil, which falls within the same buildable layer count and is effectively equivalent for the purposes of the electromagnetic design.

4.7 Adopted Electromagnetic Model

The MATLAB PDE Toolbox model is adopted as the primary electromagnetic model for the remainder of the thesis. This choice is based on modelling scope: the ANSYS model provided an independent FEM comparison, but the Student-edition mesh cap limited it to an 8-magnet subsection. The MATLAB PDE model can represent the full 1010 mm rod, including the end regions, while still reproducing the measured central field structure. It is therefore better suited for later comparison with prototype voltage and power measurements.

The modelling work in this chapter remains electromagnetic. It establishes a measured-field-consistent model of the segmented rod and uses it to support the coil geometry selection. Later chapters combine this model with the mechanical dynamics and experimental measurements to assess prototype behaviour and scaling implications.

4.8 Chapter Summary

This chapter developed the electromagnetic model used for the coil design and later thesis analysis. Since the segmented magnetic rod was not supplied with a detailed datasheet, its modelling inputs were obtained from rod identification, external measurement, surface-field scans, and interpretation of the measured field. The radial-field profile was used to infer the internal magnet-spacer segmentation: approximately 23 mm magnet length, 7 mm spacer length, 30 mm structural pitch, and 60 mm magnetic repeat period.

An analytical Biot-Savart model was used to interpret the measured field. It showed that a thick-wall interpretation of the rod did not reproduce the measured transition-region gradients as well as a thin-wall interpretation with an effective magnetic radius of approximately 13.0 mm. This thin-wall interpretation was used for the explicit ANSYS geometry, while the later MATLAB PDE model used an effective-field representation tuned to the measured surface field.

Two FEM field maps were generated. The ANSYS Maxwell model used the explicit thin-wall geometry but was limited to an 8-magnet subsection by the Student-edition mesh constraint. The MATLAB PDE Toolbox model was tuned to the same measured field and allowed the full 31-magnet rod to be solved. Both models reproduced the measured periodic field structure with similar validation error.

The calibrated field maps were used to compute flux linkage, induced voltage, and matched-load power for candidate coil outer diameters. The OD sweep used the wire and winding constraints of the prototype, so each candidate diameter corresponded to an integer number of radial winding layers. Both FEM field maps identified the same optimum coil outer diameter of approximately 48.7 mm at a winding fill factor of 0.85, corresponding to about 782 turns. A sensitivity study across fill factors from 0.55 to 0.85 showed that the recommended OD shifts by less than 0.5 mm, indicating that the result is not sensitive to the exact packing density.

The MATLAB PDE Toolbox model is adopted as the primary electromagnetic model for the remainder of the thesis because it can represent the complete 31-magnet rod, including the end regions, while still reproducing the measured central field structure. The next chapter describes the physical prototype and measurement system used to test the selected architecture.

Chapter 5

Prototype Development and Experimental Platform

5.1 Role of the Prototype in the Thesis

Chapter 4 established the electromagnetic model and coil geometry for the selected linear sliding architecture. The bench-scale prototype implements that architecture as a physical test system comprising a segmented magnetic rod, a translating three-coil carriage, added moving mass, elastic restoring ropes, and onboard electrical and motion instrumentation.

The prototype was built to measure the coupled electromechanical response of the harvester under controlled laboratory conditions, rather than to represent a final buoy-integrated package. Manual carriage motion was useful for early checks, but it was not repeatable enough for quantitative experiments. A dedicated external wave simulator was therefore built to impose controlled vertical excitation while recording both the simulator motion and the harvester measurements from each run.

Together, the prototype and wave simulator provide the experimental platform used for the prototype results in Chapter 7 and for the coupled model validation in Chapter 8.

5.2 Mechanical Implementation of the Selected Architecture

The physical prototype implements the linear sliding architecture selected in Section 3.3. Electrical energy is generated by relative motion between a segmented magnetic rod and a translating three-coil carriage. In the prototype, the rod is fixed and the carriage moves along the rod axis, carrying the coils, added masses, and onboard electronics as one moving assembly.




Figure 5.1: Bench-scale prototype assembly showing the segmented magnetic rod and translating three-coil carriage.

5.2.1 Support Structure, Carriage, and Moving Mass

The prototype is supported by two steel end plates mounted at the ends of the magnetic rod. Each plate is 2 mm thick and measures 30 cm by 16 cm. The axial plate-to-plate span is 1010 mm, equal to the rod length. In addition to supporting the assembly, the end plates provide the anchor points for the elastic restoring ropes described in Section 5.2.4.

The moving carriage consists of three stacked 3D-printed coil reels joined by printed brackets into a single translating assembly. The brackets also provide attachment points for the restoring ropes and a mounting platform for the onboard measurement electronics. Two 2.5 kg masses are attached to the carriage to increase its inertia. Including the coils, printed structure, attached masses, electronics, and mounting hardware, the moving assembly has a total mass of 6.5 kg. This is the moving mass used in the later dynamic model.

At static equilibrium measured along the vertical rod axis from the carriage ends, the carriage has approximately 66 cm of clearance to the upper end plate and 21 cm to the lower end plate. The carriage itself is approximately 14 cm long. This low equilibrium position was chosen to allow the upper restoring ropes to be as long and lightly tensioned as practical while still retaining useful carriage travel along the rod. The resulting available stroke is therefore asymmetric, with substantially more travel toward the upper end plate than toward the lower end plate.

Table 5.1: Main mechanical parameters of the bench-scale prototype.

Parameter	Value	Comment
Axial plate-to-plate span	1010 mm	Equal to the magnetic rod length
End-plate thickness	2 mm	Steel support plates
End-plate dimensions	30 cm $\times$ 16 cm	Length $\times$ width
Carriage axial length	14 cm	Total axial length of the moving carriage
Top clearance at equilibrium	66 cm	Distance from carriage top to upper end plate
Bottom clearance at equilibrium	21 cm	Distance from carriage bottom to lower end plate
Attached masses	2 $\times$ 2.5 kg	Mounted on carriage
Moving carriage assembly mass	6.5 kg	Mass used in the dynamic model

5.2.2 Prototype Coil Geometry and Assembly

The translating carriage contains three identical coils arranged axially along the magnetic rod. Each coil is wound on a 30 mm 3D-printed reel with 1.5 mm end flanges, giving an active winding length of 27 mm. As discussed in Section 4.6, this keeps the winding within one 30 mm structural pitch of the rod and avoids cancellation from neighbouring opposite-phase field regions.

The sliding clearance required over the rod, together with the 0.8 mm wall thickness of the printed reels, sets the coil inner diameter to 29.2 mm. Each prototype coil contains 800 turns, has an outer diameter of approximately 50 mm, and has an estimated DC resistance of 9.1 Ω . This outer diameter is close to the numerically preferred value identified in Chapter 4 and was retained for the physical prototype.


The prototype coils were implemented as air-core windings on 3D-printed reels. Soft-magnetic cores or flux concentrators were considered, but were not implemented because the alternating rod field would require a laminated, segmented, or otherwise low-loss core design. This was outside the available project time, so the air-core implementation was retained for model validation.

Table 5.2: Main coil parameters of the bench-scale prototype.


Parameter	Value	Unit
Number of coils	3	–
Reel axial length	30	mm
Flange thickness per side	1.5	mm
Active winding length	27	mm
Coil centre-to-centre spacing	30	mm
Turns per coil	800	–
Inner diameter	29.2	mm
Outer diameter	≈ 50	mm
Estimated DC resistance per coil	≈ 9.1	Ω

The three coils are spaced at 30 mm centre-to-centre, corresponding to one structural pitch of the magnetic rod. Since the magnetisation direction alternates between neighbouring structural sections, adjacent coils experience opposite induced-voltage polarity for the same carriage motion. The middle coil is therefore wired with reversed polarity relative to the two outer coils so that the three induced voltages add constructively. This phase relationship is retained in both the series and parallel connection configurations.

The internal arrangement of the three-coil carriage and the fabricated carriage assembly are shown in Figure 5.2.



(a) Sectional CAD view of the three-coil carriage assembly, showing the central magnetic rod, the three wound coil sections, and their axial arrangement.



(b) Fabricated three-coil carriage used in the bench-scale prototype.

Figure 5.2: Three-coil carriage geometry and fabricated prototype assembly.

5.2.3 Coil Fabrication

The prototype coils were wound in-house using the custom coil-winding rig shown in Figure 5.3. The rig was built because the harvester required three geometrically similar coils with a controlled turn count, which would have been difficult to produce consistently by manual winding alone.

The coils were wound using enamelled copper magnet wire with a diameter of 0.475 mm. The wire size was selected from the available stock as a practical compromise between winding density and electrical resistance: a finer wire would allow more turns within the same winding window but would increase resistance, while a thicker wire would reduce resistance but limit the achievable turn count.

The winder rotates the coil reel with a stepper motor while a Hall-effect sensor counts rotations from a magnet mounted on the motor axle. This allowed each coil to be wound to the prescribed 800-turn count with consistent fabrication between the three coil sections. Together with the harvester prototype and the wave simulator, the coil winder was one of the three custom hardware systems built during the project.




Figure 5.3: Custom coil-winding rig used to manufacture the prototype coils.

5.2.4 Elastic Restoring-Rope Arrangement

The final simulator-driven prototype used two elastic restoring ropes connected between the moving carriage and the upper end plate. The ropes provide the passive restoring force that returns the carriage toward its equilibrium position after displacement. Their anchor geometry is therefore part of the mechanical model used later in the thesis.

The carriage includes four possible rope-mounting holes arranged symmetrically around the rod axis. In transverse coordinates measured from the rod centreline, these holes are located at

$$(\pm 6.7 \text{ cm}, \pm 6.7 \text{ cm}). \quad (5.1)$$

In the final two-rope configuration, only the upper pair of carriage anchors is used. The corresponding upper end-plate anchors are located at

$$(12.5 \text{ cm}, 6 \text{ cm}) \quad \text{and} \quad (-12.5 \text{ cm}, 6 \text{ cm}), \quad (5.2)$$

with each end-plate anchor connected to the carriage anchor in the same quadrant.

An earlier four-rope arrangement was used during preliminary hand-driven tests, with two ropes connected toward the upper end plate and two toward the lower end plate. This was useful for initial functional checks because it constrained the carriage symmetrically. After the wave simulator was introduced, the lower ropes were removed because they did not provide a useful restoring contribution under the simulator-driven operating conditions. The final validation configuration therefore used the two upper ropes only.

The elastic ropes are nonlinear and exhibit hysteresis, so their behaviour cannot be represented accurately by a single linear spring constant. Static cyclic measurements were performed on three initial free lengths, approximately 0.82 m, 0.49 m, and 0.36 m. When expressed using engineering strain,

$$\epsilon = \frac{L - L_0}{L_0}, \quad (5.3)$$

the three datasets aligned closely enough to support a unified empirical tension–strain model.

The average rope model used for restoring-force calculations is

$$T_{\text{avg}}(\epsilon) = 97.939 \epsilon^3 - 183.753 \epsilon^2 + 132.406 \epsilon, \quad (5.4)$$

where tension is in newtons and strain is dimensionless. Separate loading and unloading fits are retained for simulations that include hysteresis explicitly, while the average curve provides the compact rope model used for first-order restoring-force calculations.




Figure 5.4: Unified empirical tension–strain model for the elastic ropes, obtained by pooling cyclic loading and unloading data from the three tested free lengths.

The measured rope law and anchor geometry provide the physical basis for the nonlinear restoring-force model developed in Chapter 6. In that model, carriage displacement changes the length and strain of each rope; the corresponding rope tension is then resolved along the rod axis to obtain the restoring force acting on the moving assembly.

5.3 Measurement and Data-Acquisition System




Figure 5.5: Measurement and data-acquisition circuit used for the bench-scale prototype.

A custom embedded measurement system was built to record the electrical output and carriage motion of the prototype during operation. The system is mounted on the moving carriage and is based on an Adafruit ESP32-S2 TFT Feather microcontroller with an ADS1115 external analog-to-digital converter. Motion sensing is provided by an AH3390Q Hall-effect sensor and an FXOS8700 magnetometer/accelerometer connected on a dedicated secondary I²C bus.

The measurement system records the signed generator voltage before path selection, rectified voltage, shunt-based current, Hall-based speed, and continuous carriage translation along the magnetic rod. The overall measurement circuit is shown in Figure 5.5, and the main hardware elements are summarised in Table 5.3.

Table 5.3: Main components of the prototype monitoring electronics.

Subsystem	Implementation	Role
Controller	Adafruit ESP32-S2 TFT Feather	Embedded measurement, local display, and web interface
ADC	ADS1115	Analog voltage measurements
Hall sensor	AH3390Q on Feather D5	Hall-edge timing, discrete position anchors, and speed estimation
Magnetometer / accelerometer	FXOS8700 on dedicated Wire1 bus	Magnetic phase tracking, acceleration measurement, and continuous translation estimation
Differential generator front-end	ADS1115 AIN0/AIN1	Signed generator voltage before path selection
Rectified voltage front-end	ADS1115 AIN2	Post-rectifier voltage measurement
Shunt current front-end	ADS1115 AIN3	Current measurement through the shunt resistor
Optional current proxy	WPSE355 to Feather A5	Auxiliary pre-switch current indication

5.3.1 Electrical Measurement Paths

The measurement circuit records both the signed generator waveform and the loaded output quantities. The generator terminal voltage is measured differentially before the path-selection switch using two ADS1115 channels biased around a mid-supply reference. If the differential ADC measurement is denoted by $V_{\text{adc,diff}}$, the generator voltage is reconstructed as

$$V_{\text{gen}} \approx 48 V_{\text{adc,diff}}. \quad (5.5)$$

This measurement preserves the polarity of the generator waveform independently of the selected load path.

The rectified output voltage is measured on a separate ADC channel through a resistive divider. If the divider output is denoted by V_{AIN2} , the rectified voltage is reconstructed as

$$V_{\text{rect}} \approx 23 V_{\text{AIN2}}. \quad (5.6)$$

Current through the external load path is measured using a low value shunt resistor,

$$R_{\text{shunt}} = 0.2 \, \Omega. \quad (5.7)$$

The corresponding current is calculated from the measured shunt voltage as

$$I_{\text{shunt}} = \frac{V_{\text{shunt}}}{R_{\text{shunt}}}. \quad (5.8)$$

The external load is replaceable, allowing the same prototype to be tested under different electrical loading conditions. A nominal $10 \, \Omega$ load was used as the default configuration, while the validation experiments later compare several load values. The firmware logs the measured voltage and current quantities and also computes derived power estimates for immediate monitoring during tests.

5.3.2 Motion Measurement and Translation Estimation

Carriage motion is measured using a Hall sensor and the FXOS8700 magnetometer/accelerometer mounted on the moving assembly. The Hall sensor provides discrete position increments and the primary speed estimates, while the magnetometer is used to refine the translation estimate between Hall events.

The Hall scaling follows from the magnetic geometry of the rod. Each magnet is 23 mm long and each spacer is 7 mm long, so one magnet-spacer section has length

$$23 \, \text{mm} + 7 \, \text{mm} = 30 \, \text{mm}. \quad (5.9)$$

The rod has a 60 mm full magnetic period, and the Hall signal changes state twice over this period. One counted Hall transition therefore corresponds to

$$\Delta x_{\text{edge}} = 30 \, \text{mm} = 0.03 \, \text{m}. \quad (5.10)$$

The Hall-based average speed is calculated from the counted edge frequency as

$$v_{\text{avg}} = f_{\text{edge}} \Delta x_{\text{edge}}, \quad (5.11)$$

where f_{edge} is the measured edge frequency. An instantaneous speed estimate can also be obtained from the time between successive edges,

$$v_{\text{inst}} = \frac{\Delta x_{\text{edge}}}{\Delta t_{\text{edge}}}. \quad (5.12)$$

To obtain continuous translation between Hall edges, the firmware uses the periodic magnetic field measured by the FXOS8700. The magnetometer components vary approximately sinusoidally along the bucking-magnet rod, and after centering and normalization the magnetic phase is estimated as

$$\theta = \text{atan2}(y_n, x_n). \quad (5.13)$$

The unwrapped phase gives a local position estimate over the 60 mm magnetic period, while Hall transitions provide trusted 30 mm anchors. The fused estimate uses the phase to interpolate within each Hall interval, with the accelerometer used for stationary detection and direction handling near reversals.

5.3.3 Embedded Interface, Sampling, and Logging


The ESP32-based firmware provides local display and browser-based access to the measured variables. Selected electrical and motion quantities can be viewed on the onboard TFT display or logged through the browser interface for offline analysis. Photographs of the measurement board are shown in Figure 5.6.

The ADS1115 is serviced using a non-blocking round-robin sequence across the differential generator-voltage, rectified-voltage, and shunt-voltage channels. With the ADS1115 configured at 860 samples/s, the effective update rate per active slot is approximately

$$\frac{860}{3} \approx 287 \text{ samples/s}. \quad (5.14)$$

The FXOS8700 is sampled separately at 400 Hz, and the firmware computes RMS, average, and speed-related quantities over a 500 ms statistics window.

During browser-based logging, the TFT display can be disabled to reduce timing load and prioritise measurement, motion estimation, and network logging.



(a) Front side of the measurement board with the microcontroller and support circuitry.

(b) Back side of the same board, showing the soldered connections.

Figure 5.6: Fabricated prototype measurement and data-acquisition electronics used during bench testing.

5.4 Integrated Prototype Assembly

Figure 5.7 shows the fully assembled moving carriage before installation in the support structure. The carriage combines the three-coil assembly, attached inertial masses, rope attachment points, and onboard measurement electronics into one translating module. This is the moving assembly represented by the 6.5 kg lumped mass in the later dynamic model.




Figure 5.7: Fully assembled moving carriage, showing the attached masses and onboard measurement electronics.

Figure 5.8 shows the same carriage installed in the full bench-scale prototype. The carriage is mounted around the segmented magnetic rod and translates along the rod axis during operation. The two attached 2.5 kg masses provide the main moving inertia, while the elastic ropes provide the passive restoring force that returns the carriage toward its equilibrium position after displacement.




Figure 5.8: Top-down view of the fully integrated bench-scale prototype, showing the segmented magnetic rod, translating three-coil carriage, onboard measurement electronics, attached masses, and elastic restoring ropes.

The integrated assembly combines the mechanical, electromagnetic, and measurement subsystems described in the preceding sections. It provides the physical test platform used for the controlled excitation and validation experiments presented later in the thesis.

5.5 Wave-Simulator System for Controlled Excitation

The prototype was first operated manually on the bench to check coil polarity, instrumentation, voltage generation, and basic carriage motion. These tests confirmed that the moving carriage produced measurable electrical output and that the embedded monitoring system recorded the relevant voltage, current, and motion signals. However, manual excitation could not provide a repeatable displacement history, so it was not suitable for quantitative comparison with the dynamic model.

A dedicated rope-driven wave simulator was therefore built to impose controlled vertical excitation on the harvester prototype. The simulator converts motor-driven wheel rotation into vertical rope displacement, allowing the prototype to be tested with prescribed motion amplitude, period, and run duration. This made it possible to repeat nominally identical test conditions while varying parameters such as excitation frequency, motion amplitude, and external electrical load.

The simulator also acts as the experiment controller for simulator-driven tests. From one browser interface, it commands the imposed motion, starts and stops paired CSV logging, and communicates with the ESP32-based generator monitor. Each run therefore produces a simulator-side record of the imposed motion and a harvester-side record of the measured electrical and carriage-response quantities.




Figure 5.9: Wave-simulator rig used to impose controlled vertical excitation during laboratory testing of the harvester prototype. The motor-driven wheel converts rotary motion into vertical rope displacement, while the counterbalance reduces the one-sided gravitational bias from the suspended load.

5.5.1 Mechanical Principle and Motion Sensing

The simulator uses a 24 V DC motor, a ZeroLAG single-channel motor driver, and a rope-driven wheel to convert rotary motion into vertical displacement. The wheel radius is 0.20 m, giving a circumference of

$$C = 2\pi R \approx 1.2566 \text{ m.} \quad (5.15)$$

One full wheel revolution therefore corresponds to approximately 1.26 m of rope travel, neglecting rope slip and the small effect of finite rope thickness.

A counterbalance is incorporated into the wheel assembly to reduce the one-sided gravitational bias from the suspended load. Without this counterbalance, the motor would require substantially different effort in the upward and downward directions. The counterbalance does not remove all asymmetry, because friction, imperfect balance, and rope geometry remain, but it makes controlled bidirectional motion around a neutral command practical.

Wheel motion is measured using 72 evenly spaced rim magnets arranged with alternating north-south polarity. In the final simulator configuration, two oppositely oriented AH3390Q unipolar Hall-effect sensors detect the alternating magnet polarities. The combined A/B Hall stream gives 72 accepted position ticks per wheel revolution. The corresponding vertical displacement per accepted tick is therefore


$$\Delta x_{\text{sim,tick}} = \frac{1.2566}{72} \approx 0.01745 \text{ m.} \quad (5.16)$$

The alternating two-sensor arrangement also provides a consistency check on the wheel odometry. During normal travel, accepted Hall events should alternate between the two sensors. Repeated activation of the same sensor can occur near reversals if the wheel moves locally in and out of one magnet's detection region without advancing to the next magnet. Such events are rejected by the simulator software and recorded in the telemetry.

The rim magnets are also observed by an FXOS8700 magnetometer mounted near the wheel. The magnetometer provides additional information on wheel direction and sub-tick motion between accepted Hall events.



(a) Unattached wheel showing the rim layout with 72 alternating magnets.



(b) Wheel mounted on the motor in the simulator test setup.

Figure 5.10: Wheel configuration used in the wave simulator.




Figure 5.11: Wave simulator circuit schematic, showing the Raspberry Pi Zero 2 W controller, motor-driver interface, 24 V motor power path, and wheel-motion sensing hardware.

5.5.2 Control Architecture and Operating Modes

The simulator is controlled by a Raspberry Pi Zero 2 W running a Flask-based backend with a browser dashboard. The Raspberry Pi sends the motor-driver PWM and enable signals, reads the Hall A/B wheel sensors, polls the FXOS8700 magnetometer over I²C, and communicates with the ESP32-based harvester monitor over HTTP during coupled experiments. Hall transitions are captured using interrupt-driven callbacks and processed by the backend to maintain the wheel odometry used for motion control and logging.

The motor driver is commanded using a neutral-centred PWM convention. A duty cycle of 50 % corresponds to neutral output, values below neutral command upward motion, and values above neutral command downward motion. The counterbalance reduces the gravitational bias from the suspended load, so the actuator can be controlled in both directions around this neutral command.

The backend separates sensing, control, supervision, and experiment logging into concurrent tasks. Hall-edge processing maintains the accepted wheel position, the magnetometer is polled at approximately 800 Hz, and the wave-control loop runs at 50 Hz. The implemented operating modes are IDLE, MANUAL, WAVE, and FAULT. Manual mode provides operator-controlled motion, while wave mode tracks a sinusoidal position reference using PID control with optional gravity and velocity feedforward terms. Idle mode holds the system near neutral when required, and fault mode latches abnormal conditions until explicitly cleared by the operator.




Figure 5.12: High-level software and signal-flow architecture of the wave simulator, showing the browser dashboard, Raspberry Pi backend, motor-control path, wheel sensing, safety supervision, experiment logging, and generator-monitor communication.

5.5.3 Supervision, Safety, and Key Parameters

The simulator includes supervision logic to limit unintended motion during web-controlled operation. Software soft limits, browser heartbeat monitoring, manual-command timeout logic, a latched kill state, and explicit fault recovery are implemented in the backend. If the browser heartbeat is lost while the rig is active, the controller enters fault mode instead of continuing motion without operator supervision. Manual commands require repeated hold updates, so manual actuation is released automatically if the command stream stops.

These safeguards were included to keep the imposed motion bounded and to ensure that loss of communication or operator input leaves the simulator in a controlled state. The main implemented parameters relevant to the simulator-driven experiments are summarised in Table 5.4.

Operating range used in the experiments. In the experimental campaign, the wave simulator was used for sinusoidal excitation over periods of approximately $T = 1.2$ s to 6 s and commanded peak amplitudes of approximately $A = 0.17$ m to 0.6 m. This range describes the tested laboratory operating envelope rather than an absolute limit of the simulator. The achieved motion was logged for each run and is used as the excitation input in the later experimental analysis and model validation.

Table 5.4: Selected implemented parameters of the wave simulator.

Parameter	Value	Comment
Wheel radius	0.20 m	Rope-driven wheel geometry
Rim magnets	72	Alternating north–south layout
Accepted Hall ticks per revolution	72	Combined A/B Hall-sensor stream
Travel per accepted Hall tick	0.017 45 m	Derived from wheel circumference
Motor supply	24 V, 17 A	Supply used for the motor driver
Motor driver	ZeroLAG 1 × 40 A	Neutral-centred PWM with enable control
PWM frequency	31 kHz	Motor-driver command frequency
Control-loop rate	50 Hz	Main wave-control update rate
Tested sinusoidal period range	1.2 s to 6 s	Range used during the experimental campaign
Tested amplitude range	0.17 m to 0.6 m	Range used during the experimental campaign
Magnetometer polling rate	800 Hz	FXOS8700 magnetic sensing
Hall speed-estimation window	0.35 s	Windowed speed estimate
Simulator log interval	0.05 s	Target simulator CSV interval
Generator fetch interval	0.01 s	Target generator-monitor polling interval
Browser heartbeat timeout	2.0 s	Fault trigger during active motion
Manual hold timeout	0.50 s	Manual-command loss protection

5.6 Coupled Simulator–Harvester Logging

During simulator-driven tests, the wave simulator also acts as the experiment coordinator. When an experiment is started from the dashboard, the system creates a shared run ID, opens a simulator-side CSV file for motion and control data, opens a harvester-side CSV file for electrical and carriage-response data, and sends the corresponding start and stop commands to the ESP32-based harvester monitor over HTTP. The two CSV files therefore form paired records of the same experimental run. The browser dashboard used for this coupled operation is shown in Figure 5.13.




Figure 5.13: Browser dashboard used for coupled simulator–harvester experiments. The left side controls the wave simulator, experiment logging, frequency sweeps, and safety limits, while the right side displays the ESP32-based generator monitor with live electrical measurements, load settings, power estimates, and carriage translation.

The logging procedure is cycle-synchronised for periodic wave tests. When the operator starts an experiment, the logger is armed and waits for the active sinusoidal reference to reach a selected phase, typically a peak or trough. The simulator and harvester recordings then begin from that repeatable point in the cycle and continue for a specified number of cycles before stopping automatically. This phase-aligned procedure improves repeatability between nominally identical runs, especially during load sweeps and frequency-response tests.

The simulator log targets a 0.05 s interval, corresponding to approximately 20 Hz. Harvester snapshots are requested at a target interval of 0.01 s, corresponding to approximately 100 Hz. These are target logging intervals rather than perfectly synchronous sample times, since the simulator and harvester monitor are separate embedded systems with independent processing and communication delays. Run coherence is instead maintained through the shared run ID, common experiment settings, phase-aligned triggering, and coordinated start–stop commands.

The simulator CSV records the commanded wave parameters, estimated simulator translation and speed, control state, and timing information. The harvester CSV records genera-

tor voltage, rectified voltage, current, load settings, and carriage-motion estimates from the onboard measurement system. This paired dataset structure provides the basis for comparing imposed excitation, measured carriage response, electrical output, and model predictions in the validation analysis.

5.7 Chapter Summary

This chapter described the physical implementation of the selected linear sliding harvester architecture. The bench-scale prototype consists of a fixed segmented magnetic rod, a translating three-coil carriage, attached inertial masses, and an elastic restoring-rope system. The moving assembly has a mass of 6.5 kg, and the measured rope tension–strain behaviour, together with the anchor geometry, provides the basis for the nonlinear restoring-force model developed in Chapter 6.

The prototype is instrumented with an onboard ESP32-S2-based measurement system using an ADS1115 for electrical measurements, a Hall sensor for discrete position and speed estimation, and an FXOS8700 magnetometer/accelerometer for continuous translation estimation. The system records the signed generator voltage, rectified voltage, load-path current, and carriage-motion quantities needed for later experimental analysis.

A dedicated rope-driven wave simulator was also developed to provide controlled vertical excitation. The simulator converts motor-driven wheel rotation into vertical rope displacement, measures the imposed motion using Hall-based wheel odometry with magnetometer assisted refinement, and tracks prescribed sinusoidal motion using a Raspberry-Pi-based control system. It also coordinates paired simulator and harvester logs through shared run IDs, phase-aligned triggering, and coordinated start–stop commands.

Together, the prototype, onboard instrumentation, and wave simulator form the experimental platform used for the experimental results in Chapter 7 and the coupled model validation in Chapter 8.

Chapter 6

System-Level Dynamic Model

6.1 Purpose and Scope

Chapter 4 established the electromagnetic field model and coil-level electrical behaviour of the selected architecture, while Chapter 5 described the mechanical prototype and measurement system. This chapter combines these elements into a system-level dynamic model for the translating carriage.

The model represents the bench-scale prototype as a one-degree-of-freedom axial oscillator. It includes the carriage mass, the nonlinear rope restoring force, mechanical damping, load-dependent electromagnetic damping, and the imposed base motion from the wave simulator. The model is used to predict the carriage response, interpret the experiments in Chapter 7, and provide the coupled model structure validated in Chapter 8.

The scope is limited to the bench-scale prototype. The model does not represent full buoy hydrodynamics, six-degree-of-freedom buoy motion, or stochastic ocean excitation directly. The validated harvester-side model is used for forward scaling in Chapter 9.

6.2 Coordinate Definition and Modelling Assumptions

The prototype is modelled as an internal carriage translating along the vertical rod axis. Let u denote the carriage position along this axis, and let u_{eq} denote the static equilibrium position. The displacement from equilibrium is defined as

$$x = u - u_{\text{eq}}. \quad (6.1)$$

The imposed motion from the wave simulator is represented as a prescribed sinusoidal base displacement,

$$Z(t) = A \sin(\omega t), \quad (6.2)$$

where A is the base-motion amplitude and ω is the excitation angular frequency. In the reference frame attached to the moving base, this appears as an inertial forcing term proportional to the base acceleration.

The model uses the following assumptions:

1. only axial carriage motion is considered,
2. the moving carriage assembly is treated as a lumped mass,
3. the elastic ropes provide the restoring force,

4. damping is represented by separate mechanical and electromagnetic contributions,
5. the prototype remains in its vertical orientation, so gravity enters through the static equilibrium force balance.

For the prototype, the moving mass is taken as

$$m = 6.5 \text{ kg.} \quad (6.3)$$

6.3 Nonlinear Restoring Force

The restoring force is provided by the elastic ropes introduced in Section 5.2.4. The rope length changes with carriage position, and the rope tension–strain relation is nonlinear. The restoring force is therefore not represented by a constant spring stiffness over the full motion range.

For each rope, the instantaneous length can be written as

$$L_i(u) = \sqrt{d_{\text{cross},i}^2 + u^2}, \quad (6.4)$$

where $d_{\text{cross},i}$ is the fixed transverse offset associated with the rope geometry. The corresponding engineering strain is

$$\epsilon_i(u) = \frac{L_i(u) - L_{0,i}}{L_{0,i}}, \quad (6.5)$$

where $L_{0,i}$ is the unloaded reference length of rope i .

Using the empirical average rope model introduced in Chapter 5, the rope tension is written as

$$T_i(\epsilon_i) = 97.939 \epsilon_i^3 - 183.753 \epsilon_i^2 + 132.406 \epsilon_i. \quad (6.6)$$

The total axial restoring force is obtained by summing the axial projections of the individual rope tensions:

$$F_{\text{rope}}(u) = \sum_i T_i(\epsilon_i(u)) \frac{u}{L_i(u)}. \quad (6.7)$$

During dynamic motion, the absolute carriage position is $u = u_{\text{eq}} + x$. The restoring-force term in the equation of motion is therefore evaluated as $F_{\text{rope}}(u_{\text{eq}} + x)$.

This restoring force is nonlinear because both the rope strain and the axial projection change with carriage position. The tangent stiffness increases with displacement amplitude, producing a hardening response. This links the rope characterisation directly to the amplitude-dependent natural frequency tested later in Chapter 7.

6.4 Damping Model

The total damping is represented as the sum of mechanical and electromagnetic contributions,

$$c_{\text{tot}} = c_{\text{mech}} + c_{\text{em}}. \quad (6.8)$$

The mechanical damping term c_{mech} represents non-electrical losses in the prototype, including rope hysteresis, friction in the moving assembly, and structural losses. As a first

approximation, these losses are treated as velocity-proportional. The corresponding damping coefficient is identified later from open-circuit free-decay measurements.

This viscous representation is an idealisation. If a significant part of the loss is caused by Coulomb friction, the equivalent viscous damping extracted from a finite-amplitude decay record will depend on oscillation amplitude. For a Coulomb-plus-viscous loss model, the equivalent damping coefficient can be written as [26]

$$c_{\text{eq}}(X) = c_{\text{visc}} + \frac{4 F_c}{\pi \omega X}, \quad (6.9)$$

where F_c is the Coulomb friction force, X is the oscillation amplitude, and ω is the natural angular frequency. A dependence of c_{eq} on $1/X$ therefore indicates a possible Coulomb component. The multi-amplitude free-decay tests in Chapter 7 are used to check this behaviour and to separate the viscous and Coulomb contributions.

The electromagnetic damping term represents the back-reaction of the generator on the carriage motion. For a linear generator, the induced electromotive force is proportional to the relative velocity, $e = K_{\text{emf}} \dot{x}$; this drives a current $i = e/R_{\text{tot}}$ through the resistive loop, producing a reaction force opposing the motion, $F_{\text{em}} = K_{\text{emf}} i$. Combining these gives $F_{\text{em}} = (K_{\text{emf}}^2/R_{\text{tot}}) \dot{x}$, so the electromagnetic damping coefficient is

$$c_{\text{em}} = \frac{K_{\text{emf}}^2}{R_{\text{tot}}}, \quad (6.10)$$

where K_{emf} is the electromechanical coupling constant and R_{tot} is the total current-loop resistance, including the equivalent coil resistance and the external circuit resistance. This damping contribution depends on the electrical configuration and therefore changes with load resistance and coil interconnection.

6.5 Equation of Motion

In the reference frame attached to the moving base, the carriage is modelled using the standard base-excited single-degree-of-freedom oscillator framework [27]. The forced carriage motion is written as

$$m\ddot{x} + c_{\text{tot}}\dot{x} + F_{\text{rope}}(u_{\text{eq}} + x) - mg = mA\omega^2 \sin(\omega t). \quad (6.11)$$

At static equilibrium, the rope force balances the carriage weight:

$$F_{\text{rope}}(u_{\text{eq}}) = mg. \quad (6.12)$$

Substituting this equilibrium balance gives

$$m\ddot{x} + c_{\text{tot}}\dot{x} + [F_{\text{rope}}(u_{\text{eq}} + x) - F_{\text{rope}}(u_{\text{eq}})] = mA\omega^2 \sin(\omega t). \quad (6.13)$$

Equation (6.13) is the governing nonlinear forced-oscillator model for the bench-scale prototype. The stiffness is displacement-dependent through the rope geometry and tension-strain relation, while the damping depends on the electrical load through the electromagnetic coupling.

6.6 Linearised Analysis

Although the full model is nonlinear, a small-amplitude linearisation around equilibrium is useful for interpreting natural frequency, damping ratio, and base-excitation response.

6.6.1 Small-Amplitude Natural Frequency

For sufficiently small displacements, the restoring-force term can be linearised as

$$F_{\text{rope}}(u_{\text{eq}} + x) - F_{\text{rope}}(u_{\text{eq}}) \approx k_{\text{eff}}x, \quad (6.14)$$

where k_{eff} is the tangent stiffness at equilibrium. The corresponding small-amplitude natural frequency is

$$\omega_n = \sqrt{\frac{k_{\text{eff}}}{m}}, \quad (6.15)$$

or equivalently

$$f_n = \frac{1}{2\pi} \sqrt{\frac{k_{\text{eff}}}{m}}. \quad (6.16)$$

This predicted natural frequency is compared with the free-decay measurements in Chapter 7.

6.6.2 Base-Excitation Transmissibility

For the linearised system, the steady-state displacement amplitude ratio is [27]

$$\frac{X}{A} = \frac{r^2}{\sqrt{(1 - r^2)^2 + (2\zeta r)^2}}, \quad (6.17)$$

where

$$r = \frac{\omega}{\omega_n} \quad (6.18)$$

is the frequency ratio and

$$\zeta = \frac{c_{\text{tot}}}{2m\omega_n} \quad (6.19)$$

is the damping ratio.

This relation shows how the carriage response depends on excitation frequency and total damping. Since the electromagnetic damping is load-dependent, changing the electrical load also changes the mechanical response of the carriage.

6.6.3 Coupled Electrical Loading and Useful Power

At the coil level, the induced voltage scales with relative carriage speed through the electromechanical coupling constant. The average power dissipated in the external load can be written as

$$P_{\text{load}} = \frac{K_{\text{emf}}^2 v_{\text{rms}}^2 R_L}{R_{\text{tot}}^2}, \quad (6.20)$$

where v_{rms} is the RMS carriage speed, R_L is the external load resistance, and R_{tot} is the total current-loop resistance.

The electrical load affects the system in two ways. It changes the electrical power division between the load and internal resistance, and it changes the electromagnetic damping applied to the carriage. The optimal load is therefore a coupled dynamic result, not only an impedance-matching result.

6.7 Nonlinear Frequency Response

The linearised model is useful for interpretation, but the full response remains nonlinear because the restoring force depends on displacement amplitude. The model is therefore also integrated numerically in the time domain over a sweep of excitation frequencies.

This nonlinear frequency-response analysis is used to predict several features relevant to the experiments:

1. resonance shifts caused by the amplitude-dependent restoring force,
2. a hardening trend from the nonlinear rope stiffness,
3. load-dependent response amplitude and power due to electromagnetic damping,
4. response limits imposed by the finite prototype stroke.

These predictions are compared with the frequency-response experiments in Chapter 7.

6.8 Model Predictions to be Tested Experimentally

The model produces two types of testable output. Some quantities follow directly from the prototype geometry and the empirical rope model. Others require damping and electromechanical parameters that must be identified from experiments.

6.8.1 Quantities Predicted from Geometry and Rope Model

Using the measured rope anchor coordinates, the calibrated rope free length of approximately 0.45 m, and the empirical tension–strain model from Chapter 5, the following quantities can be predicted before dynamic testing:

- The static equilibrium position corresponds to an equilibrium rope length of approximately 0.66 m, consistent with the value observed on the physical prototype.
- The linearised tangent stiffness at equilibrium is approximately 107 N/m, giving a predicted small-amplitude natural frequency of approximately 0.65 Hz (period ≈ 1.55 s).
- Because the restoring force has a positive stiffness gradient, the model predicts a hardening response: the effective natural frequency should increase with oscillation amplitude.
- The available stroke is asymmetric about equilibrium. At static equilibrium, the top of the carriage is approximately 66 cm below the upper plate, while the bottom of the carriage is approximately 21 cm above the lower plate. The maximum geometric travel from equilibrium is therefore about 66 cm upward and 21 cm downward, making the downward direction the limiting side for large-amplitude motion.

6.8.2 Predictions Requiring Experimentally Identified Parameters

Some predictions require parameters that cannot be obtained from geometry alone. These parameters are identified from the free-decay and wave-driven experiments in Chapter 7:

- The mechanical loss parameters c_{visc} and F_c are identified from open-circuit free-decay measurements at multiple release amplitudes. These parameters define the baseline non-electrical dissipation against which electromagnetic extraction competes. A single equivalent viscous value c_{mech} can also be fitted at a chosen amplitude, but this is only a reduced representation of the loss behaviour.
- The electromechanical coupling constant K_{emf} is identified from the voltage–speed relationship in wave-driven tests and from the difference between loaded and open-circuit decay rates. Once identified, it determines the electromagnetic damping for a given load resistance.
- The steady-state load–power curve and the coupled optimum load resistance require the mechanical-loss parameters, K_{emf} , and the excitation amplitude delivered by the wave simulator.
- The quantitative frequency-response curve, including the resonance amplitude and the frequency of maximum power extraction, requires the complete calibrated parameter set.

6.8.3 Summary of Testable Predictions

Table 6.1 summarises the main model predictions and the experiment type used to evaluate each one.

Table 6.1: Summary of model predictions and the experimental measurements used to evaluate them.

Prediction	Source	Experimental evaluation
Small-amplitude $f_n \approx 0.65$ Hz	Rope model + geometry	Multi-amplitude open-circuit free decay
f_n increases with amplitude (hardening)	Nonlinear rope model	Multi-amplitude open-circuit free decay
Open-circuit mechanical loss is approximately viscous	Damping model assumption	Coulomb–viscous decomposition of multi-amplitude decay
c_{mech} is independent of electrical load	Damping decomposition	Decay tests under different loads
c_{em} scales as $K_{\text{emf}}^2/R_{\text{tot}}$	Coupled EM model	Decay tests under different loads
Optimal R_L exceeds R_{coil}	Coupled load model	Steady-state load sweep
Series and parallel give similar maximum power	Coupled load model	Load sweep in both configurations
Resonance peak occurs near predicted f_n	Full system model	Frequency sweep
Power increases substantially near resonance	Full system model	Frequency sweep

Chapter 7 evaluates these predictions by identifying the required damping and electromechanical parameters and by documenting the measured prototype response. Chapter 8 then tests the complete coupled model against those measurements.

6.9 Chapter Summary

This chapter established a system-level dynamic model for the bench-scale internal wave energy harvester. The model represents the moving carriage as a one-degree-of-freedom oscillator with nonlinear rope restoring force, mechanical damping, load-dependent electromagnetic damping, and prescribed base excitation from the wave simulator. The resulting equation of motion links the electromagnetic analysis of Chapter 4 to the prototype experiments presented in the next chapter.

The model also shows why the prototype must be treated as a coupled electromechanical oscillator. The electrical load affects both the electrical power division and the carriage motion through electromagnetic damping. The next chapter presents the experiments and measured results used to evaluate these predictions, while Chapter 8 performs the final coupled-model validation.

Chapter 7

Experimental Setup and Results

7.1 Experimental Objectives

This chapter presents the experimental setup, test procedures, and measured response of the bench-scale harvester. The experimental campaign was designed to identify the main mechanical and electromechanical properties of the prototype. Free-decay tests were used to estimate natural frequency and damping, wave-driven tests were used to evaluate load-dependent behaviour under controlled sinusoidal excitation, and frequency-response tests were used to examine the resonance behaviour of the coupled system.

The chapter focuses on the experimental results themselves: what was measured, how the measurements were processed, and which prototype parameters were identified. The coupled-model validation against these measurements is carried out separately in Chapter 8.

7.2 General Experimental Methodology

7.2.1 General Test Protocol

Before each test, the carriage was left to settle to its equilibrium position. This settled position was used as the reference, and the translation estimate was reset to zero before recording. For wave-driven tests, the initial transient was removed and quantitative metrics were computed from the steady-state part of each recording.

Each wave-driven experiment produced paired simulator-side and harvester-side log files linked by a shared run identifier. The simulator log recorded the imposed base motion, while the harvester log recorded the carriage response and electrical output. The paired files were processed together, with the measured simulator motion treated as the input to the harvester and the harvester log treated as the measured response.

7.2.2 Instrumentation Accuracy and Known Limitations

The main instrumentation limitations were drift in the Hall/IMU-based translation estimate over long runs, limited sensitivity of the electrical measurement path at very small signal levels, finite update rates in the embedded logging system, and differences between commanded and achieved wave-simulator motion. These effects do not invalidate the experiments, but they limit how closely the measurements can be expected to match the idealised model used later for validation.

Wave-simulator operating range. During the experimental campaign, the usable operating range of the wave simulator covered periods of approximately $T = 1.2$ s to 6 s and commanded peak amplitudes of approximately $A = 0.17$ m to 0.6 m. The achieved period closely followed the commanded value, with fitted period errors below approximately 0.7% across the logged tests. The main practical limits were displacement amplitude and motion smoothness, rather than timing accuracy. The largest commanded amplitude was 0.6 m, corresponding to measured peak amplitudes of approximately 0.56 m to 0.60 m. The upper displacement range was mainly limited by the available suspension height and vertical clearance, so the tested range should be interpreted as a limitation of the laboratory setup rather than a timing limitation of the control system. The usable amplitude range also depended on the PID tuning used for each test. For this reason, the analysis uses the measured simulator displacement from each run rather than the commanded waveform.

7.3 Free-Decay Experiments

The free-decay experiments consisted of two complementary campaigns. The *load-dependent* campaign used a fixed nominal release amplitude of 20 cm and five external resistive loads to identify how the effective damping changed with electrical load. The *multi-amplitude* campaign used open-circuit releases from nominal amplitudes of 5 cm to 25 cm to examine amplitude-dependent frequency behaviour and separate the open-circuit mechanical loss into viscous and Coulomb-friction components. Together, the two campaigns identify the main damping parameters used in the system-level model.

A short open-circuit lubrication check was performed before the decay campaigns to test whether rod–carriage friction had a measurable effect on the mechanical losses. The carriage was released manually from 15 cm, 20 cm, and 25 cm before and after PTFE spray was applied to the rod. After lubrication, the equivalent damping ratio decreased on average from approximately 0.180 to 0.126, while the settling time within ± 1 cm of equilibrium increased from approximately 1.85 s to 2.80 s. Both decay campaigns were therefore performed in the lubricated condition, and residual rod–carriage friction was treated as part of the mechanical loss to be identified.

7.3.1 Purpose

The load-dependent campaign was used to quantify how electrical loading affects the decay response. At a fixed nominal release amplitude, changes in damped frequency, decay rate, and equivalent damping coefficient were compared across the tested load resistances. Because the generator was connected directly to known resistive loads in bypass mode, the experiment also tests the expected inverse-resistance relationship for electromagnetic damping.

The multi-amplitude campaign was used to examine the open-circuit mechanical behaviour of the prototype. It tested whether the nonlinear rope restoring force produced a measurable amplitude-dependent frequency trend, and whether the mechanical loss was better represented by a Coulomb-plus-viscous model than by a single viscous damping coefficient.

7.3.2 Procedure

7.3.2.1 Load-dependent tests

For the load-dependent campaign, the wave simulator was not used. The carriage was lifted manually to a nominal release height of 20 cm and released while the translation signal was recorded. All runs used the same nominal release amplitude, so differences in decay rate could be attributed primarily to the electrical load.

The generator was operated in bypass mode for all loaded runs. Five external resistive loads were tested: 2.0 Ω , 3.2 Ω , 5.0 Ω , 9.8 Ω , and 19.7 Ω . The current shunt resistor remained in series with the external load, so the external circuit resistance was

$$R_{\text{circuit}} \approx R_L + 0.2 \Omega. \quad (7.1)$$

Six decay recordings were collected for each load, and six open-circuit recordings were collected as the mechanical reference case. The tested load cases are summarised in Table 7.1.

Table 7.1: Electrical loads used in the load-dependent free-decay campaign.

Case	External load R_L [Ω]	External circuit R_{circuit} [Ω]
1	2.0	2.2
2	3.2	3.4
3	5.0	5.2
4	9.8	10.0
5	19.7	19.9

Before each run, the carriage was allowed to settle at its equilibrium position and the translation estimate was checked against the reference. The carriage was then lifted to the target release height and released as cleanly as possible, without an intentional push.

7.3.2.2 Multi-amplitude open-circuit tests

The multi-amplitude campaign also used manual release without simulator forcing. The release amplitude was varied across five nominal levels: 5 cm, 10 cm, 15 cm, 20 cm, and 25 cm. The generator terminals were left open in all runs, so no external load path was connected. Four runs were collected for each release amplitude, giving a total of 20 recordings. The achieved release amplitude was taken from the translation signal at the moment of release, rather than from the nominal target height.

7.3.3 Analysis Method

For each decay record, the translation signal was converted to displacement relative to the settled equilibrium position. The equilibrium position was estimated from the final part of the recording, after the carriage had returned close to rest. A time window around the release event was then extracted, excluding both the manual lifting phase and the final near-stationary tail.

The damped oscillation frequency was estimated from the spacing between successive extrema. For runs with multiple clear rebounds, the logarithmic decrement between successive same-sign extrema was used to estimate the damping ratio [27]:

$$\delta = \ln\left(\frac{A_k}{A_{k+1}}\right), \quad (7.2)$$

$$\zeta = \frac{\delta}{\sqrt{(2\pi)^2 + \delta^2}}. \quad (7.3)$$

For the more heavily loaded cases, the oscillation was strongly suppressed and only one weak rebound was often visible. In those cases, the damping ratio was instead obtained from a full-curve fit of the standard second-order underdamped free response [27], where the fit converged reliably:

$$\Delta x(t) = x_0 e^{-\zeta \omega_n t} \left[\cos(\omega_d t) + \frac{\zeta}{\sqrt{1 - \zeta^2}} \sin(\omega_d t) \right], \quad \omega_d = \omega_n \sqrt{1 - \zeta^2}. \quad (7.4)$$

Here x_0 is the fitted initial displacement, ω_n is the fitted natural angular frequency, and ζ is the fitted damping ratio.

The corresponding equivalent viscous damping coefficient was calculated as

$$c_{\text{tot}} = 2m\zeta\omega_n, \quad (7.5)$$

with the moving carriage mass taken as $m = 6.5$ kg.

For the load-dependent campaign, the resistance relevant to electromagnetic damping is the total current-loop resistance. This includes the equivalent coil resistance, external load, and current shunt. In the bypass-mode decay configuration, the three coils were connected in parallel. The measured parallel coil-bank resistance was

$$R_{\text{coil,par}} \approx 3.4 \, \Omega. \quad (7.6)$$

This is the directly measured terminal resistance of the parallel coil bank. It is slightly higher than one third of the individual coil resistance, which is consistent with small lead, contact, and connection contributions in the practical bypass circuit. The total current-loop resistance used for the electromagnetic damping comparison was therefore

$$R_{\text{tot}} = R_{\text{coil,par}} + R_{\text{circuit}}. \quad (7.7)$$

The open-circuit runs at the 20 cm release amplitude were used as the mechanical reference case. For each loaded case, the electromagnetic damping contribution was estimated by subtracting the open-circuit mechanical baseline:

$$c_{\text{em}} = c_{\text{tot}} - c_{\text{mech}}. \quad (7.8)$$

For an induced emf driving a resistive current loop, the expected first-order relationship is

$$c_{\text{em}} \approx \frac{K_{\text{emf}}^2}{R_{\text{tot}}}, \quad (7.9)$$

where K_{emf} is the voltage constant relating generated emf to relative velocity. The inferred c_{em} values were fitted against $1/R_{\text{tot}}$ using a through-origin fit. A separate voltage-speed check was also performed by fitting measured generator voltage against reconstructed carriage velocity during the main decay phase.

For the loaded cases, the electrical energy dissipated in the bypass path during each decay event was estimated from

$$E_{\text{bypass}} = \int p_{\text{bypass}}(t) dt. \quad (7.10)$$

This quantity is the energy dissipated in the bypass path, not stored battery energy.

For the multi-amplitude campaign, each reliable run was first reduced to one fitted $(f_n, \zeta, c_{\text{tot}})$ triplet. The amplitude dependence of the damping was then interpreted using the Coulomb-plus-viscous equivalent damping relation

$$c_{\text{eq}}(X) = c_{\text{visc}} + \frac{4F_c}{\pi\omega X}, \quad (7.11)$$

where X is the oscillation amplitude, ω is the oscillation angular frequency, c_{visc} is the velocity-proportional damping coefficient, and F_c is the Coulomb-friction force.

The preferred mechanical-loss decomposition was performed cycle-by-cycle using same-sign extrema pairs. For a pair (A_k, A_{k+2}) separated by one full oscillation cycle, the cycle log decrement, damping ratio, and equivalent damping coefficient are

$$\delta_k = \ln\left(\frac{A_k}{A_{k+2}}\right), \quad \zeta_k = \frac{\delta_k}{\sqrt{(2\pi)^2 + \delta_k^2}}, \quad c_{\text{eq},k} = 2m\zeta_k\omega_{\text{pair},k}. \quad (7.12)$$

The representative amplitude of the pair was taken as the geometric mean:

$$X_{\text{pair},k} = \sqrt{A_k A_{k+2}}. \quad (7.13)$$

Arithmetic and harmonic means were also evaluated as sensitivity cases.

7.3.4 Results: load-dependent campaign

Figure 7.1 summarises the load-dependent free-decay response. The open-circuit case shows the longest ringing and largest rebound amplitudes. As the load resistance is reduced, the oscillation is suppressed more strongly. The corresponding damping plot shows the inferred electromagnetic damping contribution against $1/R_{\text{tot}}$.




Figure 7.1: Load-dependent free-decay response. The displacement traces show the effect of electrical loading on the decay rate, while the damping plot shows the loaded cases following the expected inverse-resistance trend.

The extracted summary metrics are listed in Table 7.2. Under open-circuit conditions, the prototype remained clearly underdamped, with a mean damping ratio of 0.1040 ± 0.0113 and a mean damped oscillation frequency of 0.831 ± 0.015 Hz. Connecting a resistive

load increased the damping substantially. At the highest tested external circuit resistance, $R_{\text{circuit}} = 19.9 \Omega$, the mean damping ratio increased to 0.2311 ± 0.0045 . At the lowest tested external circuit resistance, $R_{\text{circuit}} = 2.2 \Omega$, it reached 0.7008 ± 0.0170 .

The damped frequency decreased as the effective damping increased. Across the loaded cases, the damped frequency decreased from $0.806 \pm 0.019 \text{ Hz}$ at 19.9Ω to $0.596 \pm 0.014 \text{ Hz}$ at 2.2Ω . The settling time also decreased, from $3.55 \pm 0.27 \text{ s}$ in the open-circuit case to $1.02 \pm 0.03 \text{ s}$ at the lowest tested resistance.

Table 7.2: Summary of extracted free-decay metrics for each electrical load. Reported values are mean $\pm$ standard deviation over six runs.

Case	$R_{\text{circuit}} [\Omega]$	$f_d [\text{Hz}]$	$\zeta [-]$	Settling time [s]
Open circuit	—	0.831 ± 0.015	0.1040 ± 0.0113	3.55 ± 0.27
19.7Ω	19.9	0.806 ± 0.019	0.2311 ± 0.0045	1.94 ± 0.03
9.8Ω	10.0	0.798 ± 0.015	0.2945 ± 0.0539	1.51 ± 0.03
5.0Ω	5.2	0.742 ± 0.012	0.4608 ± 0.0276	1.58 ± 0.03
3.2Ω	3.4	0.696 ± 0.011	0.5543 ± 0.0195	1.39 ± 0.02
2.0Ω	2.2	0.596 ± 0.014	0.7008 ± 0.0170	1.02 ± 0.03

The damping decomposition is summarised in Table 7.3 and shown in Figure 7.2. The open-circuit runs at the 20 cm calibration amplitude give

$$c_{\text{mech}} \approx 7.10 \pm 0.77 \text{ N s/m}.$$

Subtracting this mechanical baseline from the loaded cases gives electromagnetic damping contributions from approximately 8.68 N s/m to 40.76 N s/m. A through-origin fit of Equation (7.9) gives

$$K_{\text{emf,decay}} \approx 14.6 \text{ V s/m},$$

with $R^2 \approx 0.997$ for the loaded cases.




Figure 7.2: Free-decay damping and bypass-path energy. The inferred electromagnetic damping increases as load resistance is reduced, while the bypass-path energy is highest at an intermediate load.

Table 7.3: Damping decomposition from the load-dependent free-decay tests. The listed R_{tot} includes the measured parallel coil-bank resistance and the external circuit resistance, and c_{em} is obtained by subtracting the open-circuit mechanical baseline from the loaded-case total damping.

Case	R_{tot} [Ω]	c_{tot} [N s/m]	c_{em} [N s/m]
19.7 Ω	23.3	15.78 ± 0.30	8.68 ± 0.83
9.8 Ω	13.4	20.11 ± 3.68	13.01 ± 3.76
5.0 Ω	8.6	31.47 ± 1.88	24.36 ± 2.03
3.2 Ω	6.8	37.86 ± 1.33	30.75 ± 1.54
2.0 Ω	5.6	47.86 ± 1.16	40.76 ± 1.39

As a separate electromechanical check, the generator constant was also estimated from the voltage–speed relationship during the main decay phase. The open-circuit voltage–speed fit gave

$$K_{\text{emf,OC}} = 12.49 \text{ V s/m},$$

while the mean estimate across the usable loaded and open-circuit cases was

$$K_{\text{emf,V-v}} = 12.45 \pm 0.52 \text{ V s/m}.$$

This voltage–speed estimate is lower than the inverse-resistance estimate, but it does not depend on subtracting a mechanical damping baseline. The per-group voltage–speed estimates from the multi-amplitude campaign are also approximately amplitude-independent within the scatter. The pooled multi-amplitude voltage–speed estimate of approximately 12.7 V s/m is therefore adopted as the canonical voltage constant for the system-level model.

The loaded free-decay records also give the electrical energy dissipated in the bypass path during each decay event. Across the tested load range, the mean bypass-path energy was approximately 0.84 J to 1.40 J. The highest value occurred near the 5.0 Ω case, while the two lowest-resistance cases dissipated less energy despite producing stronger damping.

7.3.5 Results: multi-amplitude campaign

The extracted per-group summary metrics are listed in Table 7.4. Across the four reliable groups, the achieved release amplitudes were close to the nominal targets, the fitted natural frequency remained in the range 0.822 Hz to 0.842 Hz, and the voltage–speed estimate of K_{emf} was approximately amplitude-independent within the scatter. The equivalent damping coefficient c_{tot} decreased monotonically with amplitude, from 9.74 ± 0.91 N s/m at the 10 cm release to 5.81 ± 0.16 N s/m at the 25 cm release. The 5 cm group is marked as marginal because its decay records contain too few resolved extrema for a reliable second-order fit.

Table 7.4: Multi-amplitude open-circuit free-decay summary. Values are mean $\pm$ standard deviation over four runs per amplitude group. The 5 cm group is marginal because its decay records contain three or fewer extrema.

Nominal	Achieved [cm]	f_d [Hz]	f_n [Hz]	ζ [-]	c_{tot} [N s/m]	K_{emf} [V s/m]
5 cm*	4.90 ± 0.09	0.924 ± 0.017	0.963 ± 0.021	0.282 ± 0.011	22.16 ± 1.35	12.79 ± 0.31
10 cm	9.63 ± 0.44	0.834 ± 0.016	0.842 ± 0.017	0.141 ± 0.011	9.74 ± 0.91	13.17 ± 0.47
15 cm	14.51 ± 0.39	0.820 ± 0.011	0.825 ± 0.011	0.107 ± 0.002	7.20 ± 0.10	12.92 ± 0.74
20 cm	19.08 ± 0.94	0.830 ± 0.005	0.834 ± 0.006	0.093 ± 0.004	6.36 ± 0.32	11.88 ± 1.68
25 cm	22.89 ± 0.47	0.819 ± 0.006	0.822 ± 0.006	0.087 ± 0.002	5.81 ± 0.16	12.80 ± 1.18

* Marginal: excluded from quantitative fits.

Figure 7.3 shows the same trends graphically. The achieved release amplitudes follow the nominal targets closely, while the fitted natural frequency is nearly constant across the four reliable groups. A linear regression of f_n against achieved release amplitude gives a slope of -0.0012 Hz/cm, with $R^2 = 0.56$ and $p = 0.25$. The reliable data therefore do not support a statistically significant hardening trend over the 10 cm to 25 cm release-amplitude range. The voltage-speed estimate of K_{emf} is also approximately amplitude-independent, with a pooled mean close to 12.7 V s/m.




Figure 7.3: Multi-amplitude open-circuit decay summary. Top: release repeatability and fitted natural frequency. Bottom: open-circuit damping ratio and voltage-speed estimate of K_{emf} . Marginal groups are plotted with hollow markers and excluded from the regression.

The monotonic decrease in c_{tot} with release amplitude indicates that the open-circuit mechanical loss is not purely viscous. A Coulomb-friction component gives a larger equivalent damping coefficient at smaller amplitudes, as described by Equation (7.11). A group-

level fit of c_{tot} against $1/X$ across the reliable amplitude groups gives

$$c_{\text{visc}}^{(\text{grp})} = 2.89 \pm 0.24 \text{ N s/m}, \quad F_c^{(\text{grp})} = 2.67 \pm 0.14 \text{ N},$$

with $R^2 = 0.995$. This fit is retained as a conservative envelope-level sensitivity estimate.




Figure 7.4: Group-level Coulomb-plus-viscous decomposition of the open-circuit damping. (a) Equivalent damping coefficient versus release amplitude with the fitted Coulomb-viscous model from Equation (7.11). (b) The same data plotted against $1/X$, where the Coulomb component appears as a linear trend.

Cycle-resolved decomposition

The preferred mechanical-loss decomposition was obtained from the cycle-resolved analysis. Applying Equations (7.12)–(7.13) to the reliable runs produced a pooled set of same-sign extrema pairs. Figure 7.5 shows the resulting fit using the geometric-mean amplitude convention, and Table 7.5 shows the sensitivity to the pair-amplitude definition. The geometric-mean result is adopted:

$$c_{\text{visc}} = 3.40 \pm 0.15 \text{ N s/m}, \quad F_c = 0.95 \pm 0.02 \text{ N}, \quad R^2 = 0.96.$$




Figure 7.5: Cycle-resolved Coulomb-plus-viscous decomposition. Each marker represents one same-sign extrema pair from one decay run, with X_{pair} taken as the geometric mean of the two peak amplitudes.

Table 7.5: Sensitivity of the cycle-resolved decomposition to the amplitude-pair mean. The geometric-mean variant is adopted as the preferred result; the arithmetic and harmonic variants are retained as bounds.

Pair-amplitude convention	c_{visc} [N s/m]	F_c [N]	R^2
Geometric mean (preferred)	3.40 ± 0.15	0.95 ± 0.02	0.96
Arithmetic mean	2.67 ± 0.18	1.24 ± 0.03	0.95
Harmonic mean	4.04 ± 0.14	0.72 ± 0.02	0.96

At the 20 cm calibration amplitude, the group-level and cycle-resolved decompositions predict nearly the same total equivalent damping. Their main difference is the split between velocity-proportional and Coulomb loss. The cycle-resolved geometric-mean decomposition is therefore used as the mechanical-loss parameter set for the system-level ODE:

$$(c_{\text{visc}}, F_c) = (3.40 \text{ N s/m}, 0.95 \text{ N}). \quad (7.14)$$

7.3.6 Discussion

The free-decay campaigns identify three main features of the prototype response. First, electrical loading increases the effective damping as expected. The loaded decay traces settle much more quickly than the open-circuit traces, and the inferred electromagnetic damping follows the inverse-resistance trend with $R^2 \approx 0.997$. This supports the load-dependent damping form used in the system model.

Second, the open-circuit natural frequency is nearly constant over the reliable 10 cm to 25 cm release range. The reliable groups give f_n values clustered around 0.83 Hz, and no statistically significant amplitude trend is detected over this range. The static rope model predicts a small-amplitude natural frequency of approximately 0.65 Hz, which is below the finite-amplitude value measured in the reliable decay groups. The present data therefore indicate that the static rope model does not by itself reproduce the observed dynamic frequency over the tested range. The marginal 5 cm group is not used to infer the small-amplitude limit because its motion approaches the translation drift and noise floor.

Third, the open-circuit mechanical loss is better represented by a Coulomb-plus-viscous model than by a single viscous coefficient. The equivalent damping coefficient decreases with release amplitude, and the cycle-resolved decomposition gives the preferred loss parameters $c_{\text{visc}} = 3.40 \text{ N s/m}$ and $F_c = 0.95 \text{ N}$. The group-level decomposition gives a larger Coulomb term and is retained as a conservative sensitivity estimate.

Overall, the prototype behaves as a coupled electromechanical oscillator with load-dependent electromagnetic damping, a finite-amplitude natural frequency of approximately 0.83 Hz over the reliable tested range, and mechanical loss represented by both viscous and Coulomb-friction components. The parameters identified from the free-decay experiments are summarised in Table 7.6.

Table 7.6: Key parameters identified from the free-decay experiments.

Quantity	Identified value	Basis
Open-circuit damped frequency at 20 cm	0.831 ± 0.015 Hz	Load-dependent campaign
Reliable finite-amplitude natural frequency	≈ 0.83 Hz	Multi-amplitude campaign
Open-circuit equivalent damping at 20 cm	7.10 ± 0.77 N s/m	Load-dependent campaign
Adopted voltage constant K_{emf}	≈ 12.7 V s/m	Multi-amplitude voltage–speed fit
Preferred c_{visc}	3.40 ± 0.15 N s/m	Cycle-resolved decomposition
Preferred F_c	0.95 ± 0.02 N	Cycle-resolved decomposition

7.4 Steady-State Wave-Driven Experiments

7.4.1 Purpose and Analysis Strategy

The steady-state wave-driven experiments characterised the prototype under controlled periodic excitation. In contrast to the free-decay tests in Section 7.3, where the response followed from an initial displacement, the wave-driven tests imposed approximately sinusoidal base motion using the external wave simulator. This allowed the coupled mechanical and electrical response of the harvester to be measured under repeatable forcing.

The experiments address four aspects of the prototype response. The benchmark load sweep evaluates how external resistance affects useful power and carriage motion near the strongest-response region of the bench-scale system. The frequency-response experiments test whether the forced response contains a resonance peak consistent with the natural timescale identified from free decay. The harmonic-response analysis examines whether long-period excitation produces carriage motion mainly at the forcing frequency or whether higher harmonics contribute significantly. Finally, measured electrical power is compared with carriage-velocity metrics to check consistency with the electromagnetic damping model developed in Chapter 6.

All wave-driven analyses used paired wave-simulator and generator-monitor logs. The simulator log provided the measured base motion, while the generator-monitor log provided generator voltage, current estimate, bypass-path power, and carriage translation. The first two commanded cycles and the final cycle of each run were excluded from the quantitative metrics to remove startup and shutdown transients.

For the benchmark load sweep, where the forcing period was short and the response was approximately periodic at the imposed frequency, simulator and carriage amplitudes were estimated using a sinusoidal least-squares fit with a linear drift term,

$$x(t) = x_0 + x_1 t + b \sin(\omega t) + c \cos(\omega t), \quad (7.15)$$

where $x(t)$ is either the measured simulator position or the measured carriage translation. The fitted oscillation amplitude was then computed as

$$A = \sqrt{b^2 + c^2}. \quad (7.16)$$

The linear term accounts for slow drift in the carriage translation estimate, so the fitted amplitude represents the oscillatory motion rather than the absolute position estimate.

For the frequency-response and amplitude-response analyses, the same trimming and pairing procedure was used, but the motion signals were also analysed in the frequency domain. This was necessary because the long-period tests showed noticeable harmonic content. The fundamental response, higher harmonic amplitudes, peak-to-peak motion, and reconstructed RMS carriage velocity were therefore treated separately when interpreting resonance behaviour and power production.

7.4.2 Benchmark Resonance-Region Load Sweep

The first steady-state wave-driven experiment was a benchmark load sweep near the strongest response region of the bench-scale prototype. The commanded simulator motion had an amplitude of 0.17 m and a period of 1.2 s. This period was chosen because the free-decay results and preliminary forced tests indicated a strong response in this frequency range. The test was therefore used as a controlled high-signal benchmark for comparing electrical loads, not as a direct representation of full-scale ocean-wave periods.

All runs were performed in bypass mode. In this configuration, the generator was connected to an approximately resistive external load without the nonlinear behaviour of the rectified charging path. Four external load resistances were tested: 3.2 Ω , 5.0 Ω , 9.8 Ω , and 19.7 Ω . The 0.2 Ω current shunt remained in series with the bypass path, so the external bypass-path resistance was approximated as

$$R_{\text{tot}} \approx R_{\text{ext}} + 0.2 \Omega. \quad (7.17)$$

Three repeated runs were recorded for each load, giving a total of 12 runs. An open-circuit case was not included because the carriage risked impacting the lower steel plate under the same excitation condition.

The benchmark results are summarised in Table 7.7. The highest measured mean bypass-path power was obtained with the 9.8 Ω external load, corresponding to an external bypass-path resistance of approximately 10.0 Ω including the shunt. At this load, the mean bypass-path power was 4.78 W, with a standard deviation of 0.52 W across the three repeated runs. The fitted carriage amplitude was approximately 0.180 m, while the achieved simulator amplitude was approximately 0.115 m.

Table 7.7: Benchmark resonance-region load-sweep summary. Values are mean $\pm$ standard deviation across three repeated runs.

$R_{\text{ext}} [\Omega]$	$R_{\text{tot}} [\Omega]$	$A_{\text{sim}} [\text{m}]$	$A_{\text{carriage}} [\text{m}]$	$V_{\text{rms}} [\text{V}]$	$P_{\text{bypass}} [\text{W}]$
3.2	3.4	0.127 ± 0.001	0.124 ± 0.000	3.19 ± 0.09	3.07 ± 0.18
5.0	5.2	0.123 ± 0.004	0.143 ± 0.004	4.56 ± 0.18	4.09 ± 0.30
9.8	10.0	0.115 ± 0.003	0.180 ± 0.004	6.84 ± 0.38	4.78 ± 0.52
19.7	19.9	0.110 ± 0.003	0.198 ± 0.003	8.97 ± 0.47	4.12 ± 0.43

Figure 7.6 shows the power and motion trends from the load sweep. The measured bypass-path power did not increase monotonically as load resistance was reduced. Instead, it reached a maximum at the intermediate load. This is consistent with the expected trade-off in an electromagnetic harvester: low resistance increases current but also increases electromagnetic damping, which suppresses relative carriage motion and reduces generated voltage. High resistance allows larger carriage motion and higher RMS voltage, but

the current becomes too small to maximise power. The best tested compromise occurred near $10\ \Omega$ external bypass-path resistance.




Figure 7.6: Benchmark resonance-region load-sweep results. (a) Mean measured bypass-path power as a function of external bypass-path resistance. (b) Measured simulator amplitude and fitted carriage translation amplitude as functions of external bypass-path resistance. Error bars show the standard deviation across repeated runs.

The mechanical response also changed systematically with load. The fitted carriage amplitude increased from approximately 0.124 m at $3.2\ \Omega$ external load to approximately 0.198 m at $19.7\ \Omega$. This trend is consistent with weaker electromagnetic braking at higher resistance. The largest carriage motion did not produce the highest measured power, showing that maximising relative motion alone is not equivalent to maximising harvested electrical power.

The achieved simulator amplitude was lower than the nominal 0.17 m command for all tested loads and decreased slightly as the electrical resistance increased. This is consistent with coupled simulator–harvester interaction near resonance. Higher load resistance reduces electromagnetic damping, allowing larger carriage motion. The larger internal motion increases the dynamic reaction force transmitted back to the simulator, making it more difficult for the motor to maintain the commanded displacement amplitude. The measured simulator motion was therefore used in the analysis instead of assuming that the commanded amplitude was achieved exactly.

Overall, the benchmark load sweep identifies the $9.8\ \Omega$ external load, corresponding to approximately $10.0\ \Omega$ external bypass-path resistance including the shunt, as the best tested load for this resonance-region operating point. This load was therefore used as the fixed electrical load in the frequency-response experiments that follow.

7.4.3 Frequency-Response Experiments

7.4.3.1 Purpose and Procedure

The benchmark load sweep in Section 7.4.2 characterised the prototype at a single forcing period near the strongest-response region. The frequency-response experiments extended this test by varying the forcing period while keeping the generator in bypass mode at the best-performing benchmark load, $R_L = 9.8\ \Omega$.

The main purpose was to test whether the forced response contained a resonance peak near the natural frequency identified from free decay, and whether the measured bypass-

path power increased as the forcing approached this resonance. The experiments also examined whether the carriage response was dominated by the fundamental forcing component or by higher harmonic content. This distinction is important for interpreting the long-period results, where simulator tracking errors and nonlinear restoring effects can both contribute to harmonic excitation.

Two complementary campaigns were carried out. The first, referred to here as the *variable-amplitude* sweep, tested forcing periods from 1.2 s to 6.0 s. Because the wave simulator could not maintain a fixed displacement amplitude over this full period range, the commanded amplitude was adjusted at each period. The achieved base amplitude therefore ranged from approximately 12.5 cm at $T = 1.2$ s to approximately 57 cm at $T = 6.0$ s. This campaign was used primarily to locate the resonance peak.

The second campaign, referred to as the *constant-amplitude* sweep, covered $T = 3.0$ s to 6.0 s in 0.5 s steps. The commanded amplitude was 0.4 m, and the achieved base amplitude remained close to this value across the tested periods. This campaign was used to document the off-resonance falling branch under more controlled amplitude conditions. Three independent runs were analysed at each period in the variable-amplitude campaign, while one accepted paired run was analysed at each period in the constant-amplitude campaign. Each run lasted 10 commanded cycles.

7.4.3.2 Analysis Method

The paired generator-monitor and wave-simulator logs were processed using the same steady-state trimming approach as the benchmark load sweep: the first two commanded cycles and the final cycle were excluded before quantitative metrics were computed. Mean bypass-path power was then calculated from the logged bypass-path power over the remaining steady-state window.

Because the long-period response contained substantial harmonic content, the motion signals were analysed using both peak-to-peak amplitudes and Fourier amplitudes. After detrending, each simulator and carriage signal was uniformly resampled, Hann-windowed, and transformed using the FFT. The amplitudes of the first five harmonics of the commanded frequency were extracted as A_k for the simulator and X_k for the carriage, where $k = 1, \dots, 5$.

The main frequency-response metric is the fundamental transmissibility,

$$H_1 = \frac{X_1}{A_1}, \quad (7.18)$$

which isolates the carriage response at the imposed forcing frequency. The peak-to-peak transmissibility X/A was retained as a comparison metric, since differences between X/A and H_1 indicate harmonic contribution to the carriage motion. The measured fundamental transmissibility was compared with the linearised SDOF base-excitation prediction from Chapter 6,

$$H_{\text{SDOF}}(f) = \frac{r^2}{\sqrt{(1-r^2)^2 + (2\zeta r)^2}}, \quad r = \frac{f}{f_n}, \quad (7.19)$$

using $f_n = 0.83$ Hz from the free-decay measurements and $\zeta = 0.295$ from the 9.8 Ω load-dependent decay case, matching the electrical load used in the frequency-response sweep. As an independent check, the same expression was also fitted directly to the variable-amplitude H_1 data with f_n and ζ treated as free parameters.

Finally, the electromagnetic power scaling was checked by comparing measured bypass-path power with the reconstructed RMS carriage velocity,

$$v_{\text{rms},1-5f}^2 = \sum_{k=1}^5 \frac{(2\pi k f_0 X_k)^2}{2}. \quad (7.20)$$

For approximately linear electromagnetic damping at fixed electrical load, the measured power should scale approximately as $P = kv_{\text{rms}}^2$.

7.4.3.3 Results and Discussion

The main frequency-response result is shown in Figure 7.7. The largest response in the variable-amplitude campaign occurs at the shortest tested period, $T = 1.2$ s, corresponding to $f_0 = 0.833$ Hz. At this point, the measured fundamental transmissibility reaches $X_1/A_1 = 1.65$ at an achieved base amplitude of 12.5 cm. The constant-amplitude campaign documents the falling response from $T = 3$ s to $T = 6$ s.




Figure 7.7: Internal harvester frequency response in bypass mode at $R_L = 9.8 \Omega$. (a) Fundamental transmissibility X_1/A_1 versus forcing frequency for both campaigns, compared with the SDOF prediction from the decay-identified parameters and an independent SDOF fit to the variable-amplitude data. (b) Mean bypass-path power versus forcing frequency.

The SDOF prediction based on the free-decay parameters reproduces the measured trend well across the tested frequency range. Fitting Equation (7.19) directly to the variable-amplitude data gives

$$f_n = 0.823 \pm 0.003 \text{ Hz}, \quad (7.21)$$

$$\zeta = 0.306 \pm 0.002. \quad (7.22)$$

These values differ from the free-decay estimates of $f_n = 0.83$ Hz and $\zeta = 0.295$ by approximately -0.8% and $+4.1\%$, respectively. The SDOF prediction places the transmissibility

peak at $T = 1.10$ s with $X_1/A_1 = 1.78$, while the closest measured point gives $X_1/A_1 = 1.65$ at $T = 1.2$ s. This agreement between the free-decay and forced-response estimates supports the dynamic parameters used in the system-level model.

The constant-amplitude sweep is summarised in Figure 7.8. At $A \approx 0.4$ m, the mean bypass-path power decreases from approximately 0.81 W at $T = 3$ s to approximately 0.10 W at $T = 6$ s. Over the same range, the fundamental transmissibility decreases from approximately 0.20 to 0.06. This confirms that the tested periods lie on the off-resonance falling branch.




Figure 7.8: Constant-amplitude sweep summary at $A \approx 0.4$ m. (a) Achieved simulator amplitude versus period. (b) Mean bypass-path power versus period. (c) Carriage half peak-to-peak amplitude versus period. (d) Transmissibility from both the FFT fundamental and the peak-to-peak amplitude, compared with the SDOF prediction at the decay-identified parameters. The separation between the two transmissibility curves indicates harmonic contribution to the carriage motion.

Figure 7.8 also highlights a feature that is not visible from the fundamental response alone. The peak-to-peak transmissibility lies consistently above the FFT-fundamental transmissibility across the tested periods. The separation between these curves indicates that higher harmonics contribute substantially to the carriage motion. At $T = 4.5$ s, the distortion is especially clear: the peak-to-peak transmissibility is $X/A \approx 0.19$, while the fundamental transmissibility is only $X_1/A_1 \approx 0.09$, and the carriage THD reaches its maximum value in the sweep.

The long-period carriage response is therefore not purely sinusoidal. The total harmonic distortion of the carriage motion increases from approximately 51% at $T = 3$ s to values near or above 80% over most of the longer-period range, reaching 99% at $T = 4.5$ s. The mechanism is shown in Figure 7.9. Although the forcing fundamental lies below resonance at long periods, one of its higher harmonics can fall close to the natural frequency. At $T = 3.5$ s, the third harmonic is $3f_0 = 0.857$ Hz, approximately 3% above f_n . At $T = 4.5$ s, the fourth harmonic is $4f_0 = 0.889$ Hz, approximately 7% above f_n . At $T = 6.0$ s, the fifth harmonic is $5f_0 = 0.833$ Hz, essentially coincident with the decay-identified natural frequency. The harmonic closest to the natural frequency therefore shifts from lower-order

harmonics to higher-order harmonics as the forcing period is increased, while the carriage motion becomes increasingly influenced by harmonic content rather than by the forcing fundamental alone.




Figure 7.9: Harmonic-resonance mechanism in the constant-amplitude sweep. (a) Carriage harmonic content from 1f to 5f by forcing period. (b) Harmonic frequencies relative to the decay-identified natural frequency. Marker size is proportional to the carriage amplitude at that harmonic.

This harmonic-resonance behaviour is consistent with nonlinear forced response of the bench system. The free-decay tests in Section 7.3 show, however, that the static rope hardening prediction is not directly supported over the reliable tested amplitude range. Nonlinear coupling can still generate harmonic content when the imposed base motion is approximately periodic, and harmonic components that fall near the resonance band are preferentially amplified. This explains why measurable carriage motion and bypass-path power remain present even when the forcing fundamental is well below the natural frequency. The mechanism is also relevant for future operation in stochastic seas, where broadband excitation may contain components capable of exciting the harvester near its natural frequency even when the dominant wave period is much longer.

The power-velocity check is shown in Figure 7.10. The accepted runs from both campaigns collapse onto a single approximately linear trend, supporting the assumption that electromagnetic damping is approximately linear over the tested operating range. The fitted slope is

$$k = 12.10 \frac{W}{(m/s)^2}.$$

corresponding to an implied electromechanical coupling of approximately $K_{\text{emf}} = 14.34$ V s/m. This is approximately 13% higher than the adopted multi-amplitude voltage-speed value of 12.7 V s/m from Section 7.3.5. The difference is acceptable for the present model-identification purpose and is much smaller than the variation that would be expected if the electromagnetic damping were strongly nonlinear over the tested range.




Figure 7.10: Measured bypass-path power versus reconstructed RMS carriage velocity squared across the frequency-sweep campaigns. The near-linear collapse supports the linear electromagnetic damping assumption used in Chapter 6.

Overall, the frequency-response experiments establish the main dynamic timescale of the prototype, confirm the off-resonance power reduction at longer periods, and identify harmonic resonance as an important mechanism in the nonlinear rope-restored system. These results provide part of the experimental basis for the coupled validation in Chapter 8 and the scaling discussion in Chapter 9.

7.4.4 Amplitude-Response Experiment

7.4.4.1 Purpose and Procedure

The frequency-response experiments varied the forcing period, but did not isolate the effect of excitation amplitude at a fixed period. The amplitude-response experiment therefore fixed the forcing period at $T = 4$ s and varied the commanded wave-simulator amplitude from 0.2 m to 0.6 m. This period lies on the off-resonance falling branch identified in Section 7.4.3.3.

The experiment had three purposes. First, it tested whether the forced response showed an amplitude-dependent frequency shift consistent with the nonlinear rope model in Chapter 6. This provided a wave-driven comparison to the direct multi-amplitude free-decay results in Section 7.3.5. Second, it measured how bypass-path power scales with achieved base amplitude. Third, it tested whether the approximately linear electromagnetic power-velocity relationship observed in the frequency-response experiments also holds across a wider amplitude range.

Five commanded amplitudes were tested: 0.2 m, 0.3 m, 0.4 m, 0.5 m and 0.6 m. For each amplitude, three independent 10-cycle runs were recorded with the generator in bypass mode at $R_L = 9.8 \Omega$. As in the previous sections, paired wave-simulator and generator-monitor logs were recorded for each run.

7.4.4.2 Analysis Method

The analysis followed the same procedure used for the frequency-response experiments. The first two commanded cycles and the final cycle were removed, harmonic amplitudes

were extracted from the trimmed simulator and carriage motion signals, and the fundamental transmissibility $H_1 = X_1/A_1$ was computed for each run.

An apparent effective natural frequency was obtained by inverting the SDOF transmissibility relation in Equation (7.19) at each tested amplitude. This calculation treats the measured fundamental transmissibility as if it came from an equivalent linear oscillator, so the result should be interpreted as an apparent forced-response parameter rather than a direct natural-frequency measurement. During the inversion, the damping ratio was scaled with frequency as

$$\zeta_{\text{eff}}(f_n) = \zeta_{\text{ref}} \frac{f_{n,\text{ref}}}{f_n},$$

so that the underlying damping coefficient remained approximately fixed. Power-law scaling was estimated from log-log fits of bypass-path power and carriage amplitude against achieved base amplitude.

7.4.4.3 Results and Discussion

The achieved base amplitudes ranged from 0.090 m at the smallest command to 0.600 m at the largest. The simulator delivered only 45% of the requested amplitude at $A_{\text{cmd}} = 0.2$ m and 80% at $A_{\text{cmd}} = 0.3$ m. The achieved amplitude is therefore used throughout the analysis instead of the commanded amplitude.

The carriage waveform changed noticeably across the amplitude sweep, as shown in Figure 7.11. At the smallest commanded amplitude, the FFT third-harmonic amplitude $X_3 = 1.85$ cm exceeded the fundamental amplitude $X_1 = 1.28$ cm, giving a carriage total harmonic distortion of 168%. As the commanded amplitude increased, the fundamental grew while the third-harmonic amplitude saturated near 3.5 cm. The waveform therefore became more fundamental-dominated at larger amplitudes; at $A_{\text{cmd}} = 0.6$ m, the carriage THD had fallen to 58%. This is consistent with the harmonic-resonance mechanism identified in the frequency-response experiment: at $T = 4$ s, the third forcing harmonic is $3f_0 = 0.75$ Hz, about 10% below the natural frequency, so third-harmonic content in the base motion can be amplified by the prototype dynamics.




Figure 7.11: Reconstructed carriage motion over one normalised forcing cycle for each tested commanded amplitude at $T = 4$ s.

The bypass-path power scaling is shown in Figure 7.12. A power-law fit gives

$$P \propto A^{0.83}, \quad R^2 = 0.996. \quad (7.23)$$

This is well below the $P \propto A^2$ scaling expected from a linear system at fixed frequency. The measured bypass-path power increases from 0.18 W at the smallest trusted amplitude, $A = 0.24$ m, to 0.37 W at $A = 0.60$ m. A linear system with the same low-amplitude scaling would have produced approximately 1.13 W at the larger amplitude. The fitted power law for the FFT-fundamental carriage amplitude is $X_1 \propto A^{0.90}$, which is much closer to linear. The much lower power exponent therefore reflects that the harmonic content contributing to the velocity and power does not scale quadratically with base amplitude; in particular, the third-harmonic carriage response saturates rather than increasing proportionally with A .




Figure 7.12: Mean bypass-path power versus achieved base amplitude at $T = 4$ s and $R_L = 9.8 \Omega$. The fitted power law is compared with a linear-system $P \propto A^2$ reference. The shaded region marks the simulator-tracking-limited part of the sweep.

The apparent frequency shift inferred from the fundamental transmissibility is summarised in Table 7.8. The effective natural frequency estimated under the equivalent-linear SDOF assumption increases from approximately 0.72 Hz at the smallest achieved amplitude to approximately 0.78 Hz at the largest. Interpreted alone, this would be qualitatively consistent with a hardening restoring response. However, this wave-driven inference must be compared with the direct multi-amplitude open-circuit free-decay measurement in Section 7.3.5, where f_n remains close to 0.83 Hz across the reliable release-amplitude range.

The two experiments probe different parts of the response. The wave-driven amplitude sweep produces relatively small carriage amplitudes, from 3.5 cm to 9.0 cm in Table 7.8, while the free-decay experiment identifies the free response at larger carriage amplitudes of approximately 10 cm to 23 cm. In the wave-driven case, the third forcing harmonic at 0.75 Hz lies close enough to the natural frequency to influence the measured fundamental-

transmissibility interpretation. The apparent shift from 0.72 Hz to 0.78 Hz is therefore better treated as a forced-response indicator affected by harmonic resonance, rather than as direct evidence of a true amplitude-dependent natural frequency. The direct decay measurement is the cleaner test of the rope-hardening prediction, and it does not show a clear hardening trend over its reliable tested amplitude range.

Table 7.8: Amplitude-sweep summary at $T = 4$ s and $R_L = 9.8 \Omega$. The simulator delivery is the ratio of achieved to commanded base amplitude. The effective natural frequency is inferred from the fundamental transmissibility under a linear SDOF assumption.

A_{cmd} [m]	A [m]	Delivery	X [m]	X_1/A_1	$f_{n,\text{eff}}$ [Hz]	$\bar{P}_{\text{bypass}}$ [W]
0.2	0.090	45%	0.035	0.134	0.72	0.07
0.3	0.240	80%	0.057	0.105	0.81	0.18
0.4	0.363	91%	0.065	0.105	0.81	0.24
0.5	0.483	97%	0.076	0.110	0.79	0.29
0.6	0.600	100%	0.090	0.112	0.78	0.37

The amplitude sweep also supports the approximately linear electromagnetic power-velocity relation. Repeating the $P = kv_{\text{rms}}^2$ check over the 15 amplitude-sweep runs gave

$$k = 13.66 \text{ W}/(\text{m/s}),$$

corresponding to an implied electromechanical coupling of approximately $K_{\text{emf}} = 15.2 \text{ V s/m}$. This is consistent with the value obtained from the frequency-response experiments and remains within approximately 20% of the adopted multi-amplitude free-decay estimate of 12.7 V s/m .

The main conclusion from the amplitude-response experiment is that the off-resonance response is not governed by simple linear amplitude scaling. The small-amplitude regime is strongly affected by harmonic resonance, while the larger-amplitude response becomes more dominated by the forcing fundamental. The bypass-path power increases with amplitude, but much more slowly than the quadratic scaling expected for a purely linear system. At the same time, the power-velocity relationship remains approximately linear, indicating that the main nonlinearity in this test is mechanical rather than electromagnetic. The experiment also shows that the wave simulator should not be treated as an ideal displacement source at low commanded amplitudes. For this reason, the achieved simulator motion is used in the analysis rather than the commanded amplitude, so actuator tracking limitations are carried through as part of the measured excitation instead of being interpreted as harvester behaviour.

7.5 Chapter Summary

This chapter presented the experimental characterisation of the bench-scale prototype. The two free-decay campaigns identify the prototype's natural timescale, show that the open-circuit mechanical loss is better represented by a Coulomb-plus-viscous model than by a single viscous coefficient, and provide an amplitude-independent voltage-speed estimate of the electromechanical coupling $K_{\text{emf}} \approx 12.7 \text{ V s/m}$. The preferred cycle-resolved mechanical-loss decomposition gives $(c_{\text{visc}}, F_c) = (3.40 \text{ N s/m}, 0.95 \text{ N})$.

The wave-driven experiments confirm a clear resonance peak near the decay-identified natural frequency of approximately 0.83 Hz, document the harmonic-resonance mechanism that produces measurable response at long forcing periods, support an approximately linear electromagnetic damping relation across the tested amplitude and frequency range, and identify the bypass-path load that maximises useful power at the resonance-region operating point. The hardening prediction of the static rope polynomial is not supported in the reliable tested amplitude range.

Together, these results provide the experimental basis for the coupled model validation in Chapter 8 and for the scaling discussion in Chapter 9.

Chapter 8

Coupled Model Validation

8.1 Purpose and Scope

Chapter 4 defined the PDE-derived electromagnetic model of the segmented magnetic rod and translating coil assembly. Chapter 6 combined this electromagnetic conversion with the single-degree-of-freedom mechanical model, and Chapter 7 provided the prototype measurements used for validation. The key bench-scale quantities carried into this chapter are the finite-amplitude natural frequency, the load-dependent electromechanical response, and the preferred cycle-resolved mechanical-loss pair

$$(c_{\text{visc}}, F_c) = (3.40 \text{ N s/m}, 0.95 \text{ N}). \quad (8.1)$$

This chapter evaluates the coupled input-motion-to-power model against the measured prototype response. The electromagnetic subsystem is first validated independently by replaying measured carriage motion through the PDE-derived electrical model. The full coupled model is then driven by measured base motion and compared with the measured carriage response and bypass-path power.

Base motion is treated as the model input throughout this validation. The validated model therefore predicts carriage motion and electrical power from prescribed base motion, matching the input–output form used in Chapter 9.

Figure 8.1 summarises the validation workflow and its connection to the later scaling analysis.




Figure 8.1: Workflow used for the coupled-model validation and subsequent scaling analysis. The replay step evaluates the electromagnetic subsystem using measured carriage motion. The coupled validation then predicts carriage motion from measured base excitation, and the fixed harvester-side model is carried into the forward scaling analysis.

8.2 Coupled Model Used for Validation

The model evaluated in this chapter is the single-degree-of-freedom coupled electromechanical model developed in Chapter 6. The relative carriage displacement is denoted by q , positive in the carriage-motion coordinate used for the prototype model, and the prescribed base displacement is denoted by y_{base} . The equation solved during coupled validation is

$$m\ddot{q} + c_{\text{visc}}\dot{q} + F_c \tanh\left(\frac{\dot{q}}{v_s}\right) + F_{\text{em}}(q, \dot{q}) + F_{\text{rest}}(q) + m\ddot{y}_{\text{base}}(t) = 0, \quad (8.2)$$

where m is the moving mass, c_{visc} is the velocity-proportional mechanical-loss coefficient, F_c is the Coulomb-friction force, v_s is the Coulomb smoothing velocity, F_{em} is the electromagnetic reaction force, and F_{rest} is the restoring force. The electromagnetic force is evaluated from the PDE-derived electrical model defined in Chapter 4, using the coil geometry, coil interconnection, electrical load, and fixed electromagnetic scale factor from Section 8.4.

For prototype validation, $F_{\text{rest}}(q)$ is the measured two-rope restoring law from Chapter 6. The unscaled static rope model gives a local natural frequency of approximately 0.644 Hz, below the finite-amplitude value identified from the forced-response data. A fixed rope-force scale factor is therefore applied before the coupled validation:

$$\text{ropeForceScale} = \frac{m(2\pi f_n)^2}{k_{\text{static}}} = 1.634, \quad (8.3)$$

using $f_n = 0.823$ Hz and the local static rope stiffness at equilibrium. This factor sets the local linearised stiffness of the prototype model to the experimentally identified dynamic timescale. It is not fitted during the coupled validation.

8.3 Validation Strategy

Three comparisons are used:

1. **Electromagnetic replay validation:** measured carriage motion $q(t)$ is replayed through the PDE-derived electrical model, and predicted bypass-path power is compared with measured bypass-path power.
2. **Loss-partition transfer test:** mechanically identified loss partitions from the free-decay analysis are inserted into the coupled model and compared across the 52-run wave-driven dataset.
3. **Final fixed-parameter coupled validation:** the selected parameter set is frozen, the coupled model is driven by measured base excitation, and predicted carriage motion and bypass-path power are compared with the measured values.

The main validation metrics are the mean-power ratio $\bar{P}_{\text{pred}}/\bar{P}_{\text{meas}}$, the carriage-amplitude ratio $X_{\text{pred}}/X_{\text{meas}}$, the carriage-motion time-series coefficient of determination q_{R^2} , and the phase delay of the predicted fundamental carriage motion relative to the measurement. Instantaneous power-waveform R^2 is not used as a headline metric because the logged bypass-path power is a finite-window averaged quantity, whereas the model output is instantaneous electrical power. Mean bypass-path power is therefore used as the electrical validation metric.

8.4 Electromagnetic Replay Validation

8.4.1 Method

The electromagnetic replay evaluates the electrical model using measured carriage motion as the kinematic input, so the mechanical ODE is not required to predict motion. For each wave-driven run, the measured displacement is prescribed in the PDE-derived electromagnetic model. The replay uses the defined three-coil parallel configuration,

$$R_{\text{parallel,meas}} = 3.4 \, \Omega, \quad R_{\text{coil,branch}} = 10.2 \, \Omega, \quad (8.4)$$

with the prototype coil interconnection, where the middle coil polarity is reversed relative to the two outer coils. The periodic-window PDE lookup is used because the experiments operate in the central periodic region of the rod.

A run-wise magnetic-phase search is used to align the logged displacement with the PDE magnetic phase. These estimates are reduced to one global sensor-frame phase correction,

$$\Delta z_{\text{phase,sensor}} = -21.5 \, \text{mm}, \quad (8.5)$$

which is held fixed for the final replay pass. One electromagnetic scale factor is fitted on an 80% training subset and applied unchanged to the training and 20% holdout subsets:

$$\text{emScaleFactor} = 0.952. \quad (8.6)$$

The phase offset is only a sensor-frame correction for replaying measured motion. It is separate from the coupled-model coordinate convention, where the model equilibrium is defined at $q = 0$.

8.4.2 Results

Figure 8.2 and Table 8.1 summarise the replay results. With the global phase correction and training-derived scale factor fixed, the predicted mean bypass-path power agrees closely with the measured value across the 52 wave-driven runs. The training subset gives a mean prediction ratio of 1.011 and a zero-intercept slope of 0.997. The holdout subset gives a mean ratio of 1.020 and a slope of 0.954. Across all runs, the mean ratio is 1.013, the median ratio is 1.019, and the zero-intercept slope is 0.992.




Figure 8.2: Electromagnetic replay validation. Measured carriage motion is prescribed as the input to the PDE-derived electrical model.

Table 8.1: Electromagnetic replay validation metrics. One global scale factor is fitted on the training subset and applied unchanged to the holdout subset.

Subset	Runs	Mean power ratio	Median power ratio	Zero-intercept slope
Training	42	1.011	1.011	0.997
Holdout	10	1.020	1.034	0.954
All runs	52	1.013	1.019	0.992

The replay result indicates that the PDE-derived electrical model predicts mean bypass-path power accurately when the measured carriage motion is prescribed. Remaining power errors in the coupled validation are therefore more likely to originate from motion prediction, operating-regime dynamics, or mechanical-model approximations than from a large systematic error in the electromagnetic conversion model.

8.5 Loss-Partition Transfer Test

8.5.1 Method

The multi-amplitude open-circuit decay analysis in Chapter 7 produced two decompositions of the same mechanical-loss data. The group-level envelope fit gave

$$(c_{\text{visc}}, F_c)_{\text{group}} = (2.89 \text{ N s/m}, 2.67 \text{ N}), \quad (8.7)$$

whereas the cycle-resolved same-sign extrema-pair method gave the preferred pair in Equation (8.1),

$$(c_{\text{visc}}, F_c) = (3.40 \text{ N s/m}, 0.95 \text{ N}).$$

The two decompositions assign different fractions of the observed loss to viscous and Coulomb friction. Since the equivalent damping contribution of the Coulomb term depends on oscillation amplitude, the two pairs are not expected to transfer identically from free decay to wave-driven operation.

Both loss pairs were inserted into the coupled model and evaluated against the full 52-run wave-driven dataset. This comparison tests which independently identified free-decay loss partition transfers better into the forced operating regime.

8.5.2 Results

The transfer behaviour of the two loss partitions is summarised in Table 8.2. The group-level envelope partition under-predicts both power and carriage amplitude, with a median power ratio of 0.523 and a median carriage-amplitude ratio of 0.691. The cycle-resolved partition gives a median power ratio of 0.941, a median carriage-amplitude ratio of 0.785, and a median carriage-motion q_{R^2} of 0.905.

Table 8.2: Loss-partition transfer test on the full 52-run wave-driven dataset.

Decomposition method	c_{visc} [N s/m]	F_c [N]	Median P ratio	Median X ratio	Median q_{R^2}
Group-level envelope fit	2.89	2.67	0.523	0.691	0.874
Cycle-resolved partition	3.40	0.95	0.941	0.785	0.905

The cycle-resolved loss partition transfers better into the wave-driven dataset. It raises the median power ratio from 0.523 to 0.941, the median carriage-amplitude ratio from 0.691 to 0.785, and the median carriage-motion q_{R^2} from 0.874 to 0.905. The final coupled validator therefore uses the cycle-resolved pair of Equation (8.1).

8.6 Final Fixed-Parameter Coupled Validation

8.6.1 Model configuration

After the loss-partition transfer test, the final coupled validator was run with the fixed parameter set listed in Table 8.3. The electromagnetic scale factor is inherited from the replay validation, the branch coil resistance is the directly measured value, and the mechanical-loss pair is the cycle-resolved free-decay result. Each run is driven by a first-five-harmonic reconstruction of the measured simulator displacement, rather than by the commanded sinusoid. No additional parameter is fitted during the final coupled evaluation.

Table 8.3: Fixed parameter set used in the final coupled validation.

Quantity	Value	Source
Moving mass m	6.5 kg	Prototype measurement
Coil connection	Three coils in parallel	Prototype wiring
Branch coil resistance R_{coil}	10.2 Ω	Direct terminal-resistance measurement
Electromagnetic scale factor	0.952	Electromagnetic replay calibration
Viscous mechanical loss c_{visc}	3.40 N s/m	Cycle-resolved free-decay analysis
Coulomb friction F_c	0.95 N	Cycle-resolved free-decay analysis
Coulomb smoothing velocity v_s	0.004 m/s	Numerical regularisation
Target dynamic natural frequency f_n	0.823 Hz	Forced-response identification
Rope-force scale factor	1.634	Derived from $m(2\pi f_n)^2/k_{\text{static}}$
PDE mode	Periodic window	Central periodic operating region
Base excitation	First five measured harmonics	Simulator displacement reconstruction

8.6.2 Overall result

The final coupled validation result is summarised in Figure 8.3 and Table 8.4. Across all 52 wave-driven runs, the median run-wise power ratio is 0.941 and the mean run-wise power ratio is 1.088. The median carriage-amplitude ratio is 0.785, the median carriage-motion agreement is $q_{R^2} = 0.905$, and the median fundamental phase delay is -2.1° . The model captures the carriage waveform and phase reasonably well, but the response magnitude remains regime-dependent. The amplitude comparison is less scattered than the power comparison because half peak-to-peak displacement is a single kinematic quantity, whereas mean bypass-path power depends on velocity, waveform content, magnetic phase,

and electrical loading. Small motion-prediction errors are therefore amplified in the power metric, especially in the high-power resonance-region runs.




Figure 8.3: Final fixed-parameter coupled validation. The amplitude comparison evaluates the predicted carriage stroke, while the power comparison includes the combined effect of predicted motion, velocity, magnetic phase, and electrical loading.

A representative time-domain comparison is shown in Figure 8.4. The measured base motion contains harmonic content and tracking imperfections relative to an ideal sinusoid, and this measured excitation is used as the model input. The model reproduces the dominant carriage motion, phase, and higher-frequency modulation of the measured response, although small cycle-to-cycle differences remain.




Figure 8.4: Representative time-domain coupled-validation result for the A0.4_T3.5 case at $R_L = 9.8 \Omega$. The measured simulator base motion is shown together with the measured carriage displacement and the coupled-model prediction.

The corresponding generator-voltage snapshot is shown in Figure 8.5. The predicted voltage has comparable magnitude and timing to the logged `vgen_inst` signal, but the peak-level agreement is not exact. This is expected because small errors in predicted carriage position and velocity are amplified by the spatially periodic magnetic field. The main electrical validation metric remains mean bypass-path power across the full dataset.




Figure 8.5: Representative measured and predicted instantaneous generator-voltage snapshot for the same A0.4_T3.5 validation case. The model signal is sampled in the same manner as the logged `vgen_inst` channel.

Table 8.4: Overall final coupled-validation metrics across the full 52-run wave-driven dataset.

Metric	Value
Runs evaluated	52
Mean run-wise power ratio	1.088
Median run-wise power ratio	0.941
Zero-intercept power-fit slope	1.435
Median carriage-amplitude ratio	0.785
Median RMS-velocity ratio	0.967
Median carriage-motion q_{R^2}	0.905
Median fundamental phase delay	-2.1°

The regime-dependent validation metrics are given in Table 8.5. The amplitude sweep, frequency sweep, and constant-amplitude frequency sweep all have median power ratios below unity, between 0.788 and 0.932. The resonance-region load sweep gives a median power ratio of 1.773 and a median amplitude ratio of 1.104. Across all groups, the median carriage-motion q_{R^2} remains between 0.880 and 0.931. The residual error is therefore mainly in response magnitude rather than waveform shape.

Table 8.5: Final coupled-validation results by experimental group.

Experiment group	Runs	Median power ratio	Median amplitude ratio	Median q_{R^2}
Amplitude sweep	15	0.788	0.719	0.880
Frequency sweep	18	0.838	0.765	0.892
Constant-amplitude frequency sweep	7	0.932	0.760	0.908
Resonance-region load sweep	12	1.773	1.104	0.931

8.6.3 Interpretation of the Remaining Discrepancy

The replay and coupled-validation results indicate that the remaining discrepancy is mainly dynamic rather than electromagnetic. When measured carriage motion is prescribed, the PDE-derived electrical model predicts mean bypass-path power to within a few percent. The larger error appears when the model must also predict the carriage motion. In that case, the waveform and phase are generally captured well, but the response amplitude is not reproduced equally well in every operating regime.

The representative voltage comparison is consistent with this interpretation. The predicted instantaneous voltage has the correct scale and timing, but peak-level differences remain because generator voltage depends on both carriage velocity and magnetic phase. Small errors in predicted position or velocity can therefore produce visible voltage differences without requiring a separate voltage-domain calibration.

The main residual uncertainty is the mechanical response amplitude. The restoring force is based on a statically measured rope law scaled to match one experimentally identified dynamic natural frequency, so it does not guarantee the correct stiffness over all amplitudes and loading conditions. The first-five-harmonic base reconstruction also carries measurement and reconstruction uncertainty. These limitations define the confidence boundary of the present validation: the model is suitable as a forward screening tool for prescribed base excitation, but its response-amplitude bias should be considered when interpreting scaled-design predictions.

Chapter 9

Scaling Analysis and Design Implications

9.1 Scaling Workflow and Model Scope

This chapter uses the harvester-side model from Chapter 8 as a forward design-screening tool for the selected internal architecture. For each candidate design, the simulation returns the carriage displacement $q(t)$, velocity $\dot{q}(t)$, electromagnetic reaction force, load voltage and current, mean external-load power $\bar{P}_L$, peak-to-peak stroke q_{pp} , and the rod length L_{rod} required to contain the motion.

The varied design quantities are moving mass, coil count, load resistance, target natural frequency, electromagnetic coupling, friction, package length, and stroke-management settings. The stages do not all vary the same quantities. S1–S3 are imposed-sinusoid screens, O1–O2 are free-carriage ocean screens, B1–B2 introduce stroke management, and L1 retunes selected controlled-stroke candidates under fixed and variable electrical-load policies. The stage labels are used as shorthand throughout this chapter, and the workflow is summarised in Table 9.1. The results are screening estimates, not final design predictions.

In a deployed buoy, the harvester reaction force would also couple to hull hydrodynamics, mooring dynamics, wave forcing, and buoy heave response. Those effects are outside the present model.

Assumptions and design axes

Package envelope. The strict package target is a rod-length envelope of 1.8 m. This corresponds to a provisional 2 m man-portable buoy length with approximately 0.2 m reserved for end hardware, sealing, and structural attachment. It is a working design assumption, not a derived requirement.

Restoring law. The bench prototype used the measured two-rope restoring law from Chapter 6. The scaled studies do not retain this exact elastic-rope configuration, since the restoring mechanism is one of the components that must be designed for a future buoy. Each candidate instead specifies a target natural frequency, and the model converts this target into the effective restoring stiffness required at the carriage,

$$k_{\text{eff}} = m(2\pi f_{n,\text{target}})^2. \quad (9.1)$$

The restoring force is represented as

$$F_{\text{rest}}(q) = k_{\text{eff}}q + k_3q^3, \quad (9.2)$$

where $k_3 = 0$ in the linear screening cases. The cubic term is used only in the hardening sensitivity cases. Thus k_{eff} is the effective stiffness that a future restoring mechanism would need to provide at the carriage, not the stiffness of a specified physical spring.

Mechanical loss. Viscous damping and Coulomb friction are retained from the validated mechanical-loss model in Chapter 8. The design grid applies friction multipliers to the Coulomb term. These multipliers represent required reductions in mechanical loss for a future guide and bearing design, not tested hardware changes. The model does not scale mechanical loss with coil count or carriage length, so the low-friction cases should be interpreted as design targets.

Electrical load. The reported power $\bar{P}_L$ is the mean power delivered to the external load. Load resistance is a design variable in the screens. The circuit calculation includes the measured branch resistance and a small series contact resistance of $0.2\ \Omega$. In the later load-tuning study, the external load is either fixed, selected by representative sea state, or varied during the time-domain simulation.

Electromagnetic coupling. An electromagnetic multiplier m_{EM} scales the effective flux-linkage gradient $d\lambda/dq$, equivalently the EMF constant $K(q)$. A value above unity increases both the generated voltage and the electromagnetic reaction force. For a given load path, the induced voltage scales approximately with K , while the electromagnetic damping contribution scales approximately with K^2/R_{eq} . The higher-coupling cases require load retuning because m_{EM} changes both the electrical output and the mechanical operating point. Physically, $m_{\text{EM}} > 1$ represents an unmeasured coupling improvement, for example through stronger magnetic loading, reduced air gap, improved magnetic-circuit geometry, laminated or segmented soft-magnetic cores around the coils, or revised winding.

Base motion. The sinusoidal screens use

$$y_{\text{base}}(t) = A \sin(2\pi t/T). \quad (9.3)$$

The ocean screens use wave-follower forcing,

$$y_{\text{base}}(t) = \eta(t), \quad (9.4)$$

where $\eta(t)$ is a simulated sea-surface elevation. This assigns the full wave elevation to the harvester base and is therefore an upper-bound forcing assumption for a real buoy. A separate B1 diagnostic uses a second-order cylinder heave RAO to test sensitivity to this base-motion assumption.

Feasibility metrics. The target is $\bar{P}_L = 10\ \text{W}$. Free-carriage screens are mainly filtered by required rod length. Stroke-management and load-tuning screens also evaluate peak relative carriage acceleration, bottom-out, peak stop force, stop-contact fraction, and coil current. These thresholds are screening limits, not final survivability specifications.

Table 9.1: Summary of the scaling stages used in this chapter.

Stage	Forcing or model basis	Role in the scaling analysis
S1	Sinusoid, $A = 0.25$ m, $T = 4$ s	Initial linked-mass screen for coil count and target natural frequency.
S2	Sinusoid, $A = 0.25$ m, $T = 4$ s	Independent moving mass, coil count, load, friction, and electromagnetic coupling.
S3	Sinusoids with $T = 6$ s and 8 s	Checks the sinusoidal trends at longer imposed periods.
O1	Irregular wave-follower ocean states	Broad ocean-weighted screen before strict package filtering.
O2	Irregular wave-follower ocean states	Feasibility-first free-carriage screen with the 1.8 m envelope as a primary filter.
B1	Cylinder-heave RAO with hardening and soft-stop diagnostics	Sensitivity check for base-motion model and stroke-management force laws.
B2	Irregular wave-follower ocean states with bumper dynamics	Integrated controlled-stroke ocean screen before final load retuning.
L1	Selected B2 candidates with fixed and variable load policies	Load retuning of controlled-stroke candidates.

9.2 Sinusoidal Screening

The sinusoidal screens isolate design trends before applying irregular ocean forcing. They use the validated electrical model from Chapter 8, including the parallel coil configuration, measured branch resistance, PDE-derived electromagnetic lookup, and replay-derived electromagnetic scale factor.

Sweep S1 used $A = 0.25$ m, $T = 4$ s, $N = 3$ – 30 coils, and $f_n = 0.15$ – 0.85 Hz. Moving mass was linked to coil count by the assumed carriage-mass relation. The strongest feasible point produced $\bar{P}_L \approx 2.5$ W at $N = 30$, $m = 20$ kg, $f_n = 0.25$ Hz, and $L_{\text{rod}} = 1.77$ m. Increasing coil count alone did not approach the 10 W target.

Sweep S2 made moving mass an independent design axis. With baseline electromagnetic coupling and measured losses, the best result was approximately 6.8 W at $N = 24$, $m = 40$ kg, $f_n = 0.25$ Hz, and $L_{\text{rod}} = 1.78$ m. Added moving mass produced the largest gain over S1. Increasing coil count helped only up to a point; while it increases raw generation capacity, the added coils also increase internal resistance, carriage length, and electromagnetic damping. Above roughly 18–24 coils, the added active length no longer translated into a proportional power increase under the same package constraint.

The strongest S2 long rerun used only $N = 6$ coils with $m = 40$ kg, $f_n = 0.25$ Hz, $R_L = 24.6$ Ω , and $m_{\text{EM}} = 1.5$. It produced $\bar{P}_L \approx 10.2$ W at $L_{\text{rod}} = 1.76$ m. This does not mean that fewer coils are generally better. Rather, it indicates that for this operating point, a shorter carriage allowed for a longer usable stroke, and consequently higher velocities, while the higher electromagnetic multiplier and retuned load maintained sufficient electrical coupling.

Sweep S3 used longer periods: $T = 6$ s with $A = 0.375$ m, and $T = 8$ s with $A = 0.500$ m. The best 6 s rerun produced approximately 4.5 W; the best 8 s rerun was lower. Although the displacement amplitude was increased, the imposed base acceleration scales with $A(2\pi/T)^2$, so the longer-period cases provided less acceleration to drive the internal oscillator.

Table 9.2: Headline sinusoidal screening results. Power values are rounded to one decimal.

Result	A [m]	T [s]	$\bar{P}_L$ [W]	N	m [kg]	f_n [Hz]	k_{eff} [N/m]	R_L [Ω]	EM $\times$	L_{rod} [m]
S1 best	0.250	4	2.5	30	20	0.25	49	12.1	1.0	1.77
S2 baseline	0.250	4	6.8	24	40	0.25	99	6.0	1.0	1.78
S2 long rerun	0.250	4	10.2	6	40	0.25	99	24.6	1.5	1.76
S3, 6 s	0.375	6	4.5	9	40	0.17	44	21.7	1.5	1.74
S3, 8 s	0.500	8	2.5	9	40	0.13	25	21.7	1.5	1.73

S1 links moving mass to coil count. S2 uses independent mass selection. The S2 long rerun uses a shorter carriage and retuned load. The S3 cases have lower forcing acceleration at longer periods.

The sinusoidal screens identify three important trends. First, the strongest candidates sit at the upper end of the tested moving-mass range. Second, the best coil count is not necessarily the largest coil count, because carriage length and resistance also change. Third, the only 10 W sinusoidal point depends on improved electromagnetic coupling and an idealised 4 s imposed motion.




Figure 9.1: S1 and S2 screening under the 4 s sinusoid. (a) S1 design-space map. (b) S2 power versus coil count at $f_n = 0.25$ Hz, grouped by moving mass.




Figure 9.2: S2 and S3 long-rerun results in the $(L_{\text{rod}}, \bar{P}_L)$ plane. The highest-power sinusoidal point is the S2 4 s rerun with $m_{\text{EM}} = 1.5$. Several other S2 candidates remain near 9 W to 10 W, but the longer-period S3 cases stay below 5 W.

9.3 Sea-State Representation and Ocean-Weighted Power

The ocean screens use wave data from the Surtseyjardufl station, an in-situ observation point off the south coast of Iceland included in the Copernicus Marine Arctic Ocean in-situ product [28]. The record was quality-filtered to 120,516 (H_s, T_z) samples and binned in the H_s – T_z plane. The site is treated as a representative Icelandic coastal input for this screening study, not as a universal Icelandic wave climate.

The distribution was reduced to the ten representative states in Table 9.3. Six operational states with $H_s \leq 4$ m carry 86.4% of the normalised weight. Four higher- H_s stress states carry the remaining 13.6%.

Table 9.3: Representative sea states from the Surtseyjardufl record. Weights are normalised over all ten states.

State	H_s [m]	T_z [s]	T_p [s]	Weight [%]	Class
SS01	1.31	5.09	6.54	25.3	operational
SS02	1.40	6.81	8.75	15.9	operational
SS03	2.56	5.43	6.99	7.6	operational
SS04	2.85	6.50	8.36	14.6	operational
SS05	2.89	7.66	9.86	13.9	operational
SS06	3.92	8.89	11.43	9.1	operational
SS07	4.68	7.35	9.45	4.8	stress
SS08	5.82	8.49	10.91	3.1	stress
SS09	6.28	9.74	12.53	4.0	stress
SS10	8.68	10.81	13.90	1.6	stress

For each representative state, wave-follower base motion was generated by direct random-phase harmonic summation [29] from a Bretschneider one-sided spectrum [30],

$$y_{\text{base}}(t) = \sum_{k=1}^{N_f} \sqrt{2S(f_k) \Delta f} \cos(2\pi f_k t + \phi_k), \quad (9.5)$$

with independent uniformly distributed phases ϕ_k . The spectrum used H_s and $T_p = 1.286T_z$.

For O2, B2, and the fixed-load ocean evaluations, the synthesis used $N_f = 900$ components over 0.025 Hz to 0.50 Hz, sampled at 200 Hz, with a 200 s spinup and a 400 s evaluation window. The L1 study used $N_f = 420$ components over 0.025 Hz to 1.20 Hz, sampled at 166.7 Hz, with a 96 s spinup and a 192 s evaluation window. O1 and O2 screening runs used fewer seeds per state than the later controlled-stroke evaluations. B2, the main L1 comparison, and the strict 1.8 m confirmation used five independent seeds per representative state.

Ocean-weighted power means

$$\bar{P}_{L,\text{ocean}} = \sum_{i=1}^{10} w_i \bar{P}_{L,i}, \quad (9.6)$$

where w_i is the normalised occurrence weight of representative state i , and $\bar{P}_{L,i}$ is the mean external-load power in that state. This metric estimates expected power over the representative distribution. It does not preserve storm sequence, calm periods, or wave-group statistics.




Figure 9.3: Quality-filtered Surtseyjardufll (H_s , T_z) distribution and the ten representative states. Marker area scales with normalised occurrence weight.

9.4 Wave-Follower Ocean Screening

Sweep O1 applied the S2 design space to the representative irregular-wave states. Candidates were first ranked by a spectral power estimate and then checked with nonlinear time-domain simulation. The spectral estimate is a linear SDOF frequency-domain ranking metric that approximates candidate power from the sea-state spectrum before the shortlisted cases are evaluated with the nonlinear time-domain model.

The highest raw ocean-weighted power was approximately 15.2 W, but the corresponding required rod length was 9.34 m. The shortest-package O1 nonlinear candidate still required 5.48 m and produced approximately 5.1 W. The high raw power values therefore came from long-stroke candidates that are incompatible with the intended package.

Sweep O2 made package length a primary filter. The lower bound on f_n was raised to suppress long-stroke candidates, load resistance was treated as a fixed hardware variable, and candidates were ranked after the nonlinear stroke check. The best strict 1.8 m candidate, OC0134, produced $\bar{P}_L \approx 3.8$ W over the ten states, with $N = 6$, $m = 40$ kg, $f_n = 0.50$ Hz, $k_{\text{eff}} \approx 395$ N/m, $R_L = 8 \Omega$, $m_{\text{EM}} = 1.5$, friction multiplier 0.25, and $L_{\text{rod}} = 1.78$ m.

The relaxed-envelope candidate OC0081 produced approximately 7.5 W raw weighted power at $L_{\text{rod}} = 2.56$ m. Applying the strict 1.8 m cutout to the same time histories reduced usable power to approximately 3.9 W. The difference shows that the extra power is obtained mainly in states where the free stroke exceeds the strict package.

Table 9.4: Headline O1 and O2 results. Power values are rounded to one decimal.

Result	ID	$\bar{P}_L$ [W]	L_{rod} [m]	N	m [kg]	f_n [Hz]	k_{eff} [N/m]	R_L [Ω]	EM $\times$	F_c $\times$
O1 best power-first	OC041	15.2	9.34	9	40.0	0.17	44	21.7	1.5	0.25
O1 shortest nonlinear candidate	OC034	5.1	5.48	27	18.5	0.25	46	12.1	1.0	1.00
O2 best strict stroke-feasible	OC0134	3.8	1.78	6	40.0	0.50	395	8.0	1.5	0.25
O2 relaxed-envelope candidate	OC0081	7.5	2.56	9	40.0	0.40	253	8.0	1.5	0.25

The table shows why the ocean screens differ from the sinusoidal cases. High power remains possible in the model, but the strongest O1 candidates use low target natural frequencies and low effective stiffness, which leads to several metres of required stroke. O2 raises f_n and k_{eff} to reduce stroke toward the 1.8 m package constraint, but the strict-envelope free-carriage power falls to about 4 W. This indicates that, under the strict package constraint, directly targeting the lower dominant wave-frequency range is not necessarily beneficial: the reduced stiffness increases relative motion, but the resulting stroke demand exceeds the available envelope. The strongest O1 and O2 rows also rely on improved electromagnetic coupling and reduced Coulomb friction.




Figure 9.4: O1 wave-follower ocean screen: ocean-weighted external-load power against worst-case required rod length. The strongest candidates exceed the 1.8 m envelope.




Figure 9.5: O2 feasibility-first ocean screen. (a) Raw and 1.8 m-cutout power per final candidate. (b) Power against worst-case required rod length.

9.5 Stroke Management

The free-carriage screens identify stroke length as the main obstacle. B1 and B2 therefore test modelled stroke-management forces.

Here, a bumper is an abstract end-of-travel force that becomes active when $|q|$ exceeds a stop position q_{stop} . In the integrated simulations the stop force is

$$F_{\text{stop}}(q, \dot{q}) = k_{\text{stop}} \max(0, |q| - q_{\text{stop}})^{n_{\text{stop}}} \text{sign}(q) + c_{\text{stop}} \max(\dot{q} \text{sign}(q), 0) \text{sign}(q). \quad (9.7)$$

The first term is a polynomial spring penalty. The second term is approach-only damping, active only while the carriage moves further into the stop. The hard travel limit is the package-derived motion limit. Bottom-out means that the simulated bumper compression exceeds the allowed compression margin. Stop-contact fraction is the fraction of the evaluation window spent in bumper contact.

The stop-force, acceleration, contact-fraction, and bottom-out checks measure different failure modes. Peak stop force is the peak force from the bumper term alone. The sweep uses a two-tier force limit: a 4000 N hard ceiling used for feasibility screening, and a 2000 N design target used to set bumper stiffness. A candidate can therefore pass the stop-force check while exceeding the design target in the most energetic stress-state runs. Peak acceleration is the peak relative acceleration of the carriage coordinate and includes the combined effect of base excitation, restoring force, electromagnetic reaction, friction, and any bumper contact. A case can therefore pass the stop-force limit while still failing the acceleration limit. The 6 g acceleration threshold is used here as a provisional screening limit rather than a derived survivability requirement. It provides a practical way to reject candidates with severe impact-like carriage dynamics, but a deployable design would require this limit to be replaced by a criterion derived from guide-bearing loads, carriage structure, end-stop design, sensor mounting, and fatigue requirements.

B1 sensitivity diagnostic

B1 is a sensitivity diagnostic, not a final stroke-management design. It uses a second-order cylinder heave response-amplitude operator (RAO) instead of wave-follower base motion. The RAO maps wave elevation to buoy heave response as a function of frequency. B1 also uses a slightly different sea-state partition with maximum $H_s \approx 5.5$ m. Its absolute power values are therefore not directly comparable to O1, O2, B2, or L1. Its purpose is to show the power-package trend produced by added nonlinear restoring force.

Without hardening or bumpers, the strongest B1 candidates required 11 m to 12 m of package length and produced roughly 41 W. Cubic hardening and soft-stop bumpers reduced package length at the cost of reduced power. Hard cubic hardening spans approximately 7.2 W at 2.71 m to 11.4 W at 3.33 m. Medium cubic hardening spans approximately 10.4 W at 3.13 m to 15.4 W at 3.94 m. The hard soft-stop case reaches approximately 8.2 W at 2.5 m. These values indicate a useful trade-off direction, but B1 is not a final stroke-management design.

B2 integrated bumper sweep

The load values inherited by B2 came from the free-carriage screens, where the bumper was not yet active. Once the stop force is included in the time integration, the electrical operating point changes. The B2 results therefore mainly identify controlled-stroke geometries and failure modes; the effect of electrical retuning is evaluated in L1.

B2 includes the bumper force during time integration and evaluates final candidates over ten representative states with five seeds per state. The grid spans moving masses of 10 kg to 40 kg, coil counts of 3–15, target natural frequencies of 0.30 Hz to 0.70 Hz, load values of 1.5 Ω to 150 Ω , package classes of 1.8 m to 2.75 m, three friction multipliers of 0.25, 0.50, and 1.0 applied to the Coulomb coefficient, and bumper settings for compression margin, design stop force, and stop exponent.

No final B2 candidate passes all configured controlled-stroke checks. The strongest final candidate, LR00183, produced approximately 4.2 W at $L_{rod} = 2.51$ m, but reached

approximately 28 g peak relative carriage acceleration. The best strict-package candidate, LR00001, produced approximately 2.4 W at $L_{\text{rod}} = 1.80$ m and also failed on acceleration and bottom-out. Across the 25 final candidates, coil current, stop-force, and contact-fraction checks passed (the stop-force check is against the 4000 N hard limit; peak forces in the most energetic stress-state runs reached approximately 2550–2760 N, above the 2000 N design target), but none passed the acceleration check and 22 bottomed out in at least one state-seed combination.

All 25 final candidates belong to reduced-friction scenarios: 15 apply a Coulomb multiplier of 0.25 and 10 apply a multiplier of 0.50; no baseline-friction candidate survived to the final evaluation stage. Peak acceleration is tightly clustered at 27–28 g across all 25 candidates regardless of package class or bumper-stiffness setting, which suggests the impact acceleration is set primarily by the carriage mass and wave-loading conditions rather than by the bumper parameters explored here.

Table 9.5: Stroke-management results. Values are rounded summary results. B1 uses cylinder-RAO base motion; B2 uses wave-follower base motion with integrated bumper dynamics.

Case	$\bar{P}_L$ [W]	L_{rod} [m]	Peak accel. [g]	Note
B1 free stroke, no management	41.0	11–12	—	diagnostic only
B1 medium cubic, low-package end	10.4	3.13	—	restoring-force sensitivity
B1 medium cubic, high-power end	15.4	3.94	—	restoring-force sensitivity
B1 hard cubic, low-package end	7.2	2.71	—	restoring-force sensitivity
B1 hard cubic, high-power end	11.4	3.33	—	restoring-force sensitivity
B1 hard soft-stop bumper	8.2	2.5	—	stop-force sensitivity
B2 strict-package LR00001	2.4	1.80	28	fails accel. and bottom-out checks
B2 overall LR00183	4.2	2.51	28	fails accel. and bottom-out checks

The headline controlled-stroke candidates use package-derived bumper geometries, summarised in Table 9.6. The hard limit q_{hard} is set by the available package length after subtracting carriage length and end margins. The stop position is then placed inside this limit by the allowed bumper compression,

$$q_{\text{stop}} = q_{\text{hard}} - x_{\text{bump,max}}, \quad (9.8)$$

and the stop stiffness coefficient is set from the design force at maximum compression,

$$k_{\text{stop}} = \frac{F_{\text{design}}}{x_{\text{bump,max}}}. \quad (9.9)$$

The table includes the B2 headline rows and the later L1 headline rows, because L1 retunes electrical loading on selected B2 geometries rather than introducing a new bumper design sweep.

Table 9.6: Bumper settings for the headline controlled-stroke candidates.

ID	Role	Pkg [m]	q_{hard} [m]	$x_{\text{bump,max}}$ [m]	q_{stop} [m]	F_{design} [N]	n_{stop} [-]	k_{stop} [N/m ^{n}]	c_{stop} [Ns/m]
LR00001	strict-package B2 and best 1.8 m L1 row	1.80	0.645	0.075	0.570	2000	2	3.56×10^5	438
LR00183	strongest B2 row	2.50	0.995	0.100	0.895	2000	2	2.00×10^5	379
LR00241	strongest L1 row	2.75	1.120	0.100	1.020	2000	1	2.00×10^4	268
LR00197	second 10 W-class L1 row	2.50	0.995	0.100	0.895	2000	1	2.00×10^4	268
LR00169	strongest baseline-EM L1 row	2.50	0.995	0.150	0.845	1000	1	6.67×10^3	155

The unit of k_{stop} depends on n_{stop} , since the spring part of the stop law is $F = k_{\text{stop}}x^{n_{\text{stop}}}$. For $n_{\text{stop}} = 1$, k_{stop} has units of N/m; for $n_{\text{stop}} = 2$, it has units of N/m².

Table 9.6 shows that the final controlled-stroke candidates are not only distinguished by package length and electrical loading. They also differ in stop location, allowable compression, stop exponent, and damping. The strict 1.8 m candidate has the smallest travel allowance and uses a quadratic, high-stiffness bumper. The later 2.5 m and 2.75 m L1 headline rows generally use longer travel and mostly linear stop laws, which reduce the abruptness of the stop interaction. This is consistent with the final feasibility pattern: the limiting issue is not coil current or peak stop force, but the combination of available stroke, bottom-out avoidance, and peak acceleration.

B1 indicates that nonlinear stroke-management forces can reduce package length while preserving several watts of power. B2 shows that the bumper formulation used here does not yet control the carriage acceptably under unrefined loading. The remaining requirement is a force–displacement and damping law that keeps the carriage inside the package while limiting peak acceleration and avoiding repeated hard bottom-out; load retuning on the same geometries is treated next.


B1 stroke-management diagnostic (cylinder 2nd-order RAO forcing, different SS partition than O1/O2/B2/L1)

Figure 9.6: B1 stroke-management diagnostic under cylinder-RAO base motion. (a) Cubic hardening. (b) Soft-stop bumpers. Harder force laws reduce package length and recovered power.




Figure 9.7: B2 integrated bumper sweep. (a) Ocean-weighted power against peak relative carriage acceleration. (b) Best B2 candidate per package class.

9.6 Electrical Load Tuning

L1 retunes the electrical load for selected B2 geometries. The motivation is that the B2 candidates used load values inherited from earlier screening, while the active bumper dynamics changed the operating point. Electrical load affects both power extraction and electromagnetic damping, so retuning it can change both harvested power and stroke severity.

Three load classes were tested. A global fixed load uses one external resistance for all sea states. A scheduled fixed load uses one fixed resistance per representative sea state. A within-wave dynamic policy updates $R_L(t, q, \dot{q})$ during the time-domain simulation. The policy names in the result tables have the following meaning:

- `best_global_fixed`: one fixed resistance selected for the full ten-state weighting;
- `scheduled_fixed`: one fixed resistance selected separately for each representative state;
- `velocity_smooth`: resistance is reduced smoothly at higher carriage speed;
- `stroke_brake`: resistance is reduced as the carriage approaches the travel limit;
- `combined`: combines the velocity-based and stroke-based reductions;
- `soft_bangbang`: switches smoothly between low and high resistance based on velocity.

The suffix `_hint` means that the policy uses the candidate load hint as its base resistance. The suffix `_scheduled` means that the policy uses sea-state scheduled base resistance.

The L1 sweep used selected B2 candidates, electromagnetic multipliers of $\times 1.0$, $\times 1.2$, and $\times 1.4$, package overrides from 1.5 m to 2.75 m, and final comparisons with five seeds per representative state. It is not a new structural grid over all masses, coil counts, and natural frequencies; it is a load-policy retest of selected B2 geometries.

Headline candidates

All headline L1 candidates require reduced Coulomb friction; the B2 selection produced no passing candidates at the prototype friction level.

The strongest L1 candidate is LR00241: $N = 12$, $m = 40$ kg, $f_n = 0.45$ Hz, $k_{\text{eff}} \approx 320$ N/m, $m_{\text{EM}} = 1.4$, $F_c \times 0.25$, a combined_hint within-wave load policy with base resistance $R_L = 85 \Omega$, and a design stop force of 2000 N. It produced $\bar{P}_L \approx 10.4$ W, with seed-to-seed standard deviation 0.6 W, peak acceleration of approximately 4.9 g, and $L_{\text{rod}} = 2.69$ m inside the 2.75 m package class.

At the same package class and EM multiplier, static-load alternatives produced similar raw power but did not give a better final controlled result. The global fixed-load row reached approximately 10.6 W in the screen but bottomed out in at least one stress-state seed. The scheduled fixed-load row reached approximately 10.5 W, but exceeded the 6 g acceleration limit and bottomed out in one stress-state seed. The within-wave policy retained nearly the same power while satisfying the configured checks.

A second 10 W row is LR00197 at the 2.5 m package class, from the improved-guidance scenario ($F_c \times 0.50$). With $m_{\text{EM}} = 1.4$ and a global fixed load of 15Ω , it produced approximately 10.0 W, with seed-to-seed standard deviation 0.7 W, peak acceleration of approximately 5.6 g, and $L_{\text{rod}} = 2.45$ m. At $m_{\text{EM}} = 1.0$ ($F_c \times 0.25$), the strongest controlled L1 row at the 2.5 m class is LR00169. It uses a global fixed load of 6.8Ω and produces approximately 9.8 W, with seed-to-seed standard deviation 0.9 W, and peak acceleration of approximately 3.0 g. This is a fixed-load result at unimproved electromagnetic coupling but still requires 75% friction reduction from the prototype value; it is much higher than the corresponding B2 result because the L1 load was retuned after the bumper dynamics were included.

Table 9.7: Headline L1 results. The $\pm$ value is the seed-to-seed standard deviation of the ocean-weighted total.

Result	ID	Pkg [m]	EM $\times$	F_c $\times$	Policy	R_L [Ω]	$\bar{P}_L$ [W]	L_{rod} [m]	Peak accel. [g]
Best overall	LR00241	2.75	1.4	0.25	combined_hint	85	10.4 ± 0.6	2.69	4.9
Second ≥ 10 W row	LR00197	2.50	1.4	0.50	best_global_fixed	15	10.0 ± 0.7	2.45	5.6
Best $m_{\text{EM}} = 1.0$ row	LR00169	2.50	1.0	0.25	best_global_fixed	6.8	9.8 ± 0.9	2.43	3.0
Best strict 1.8 m	LR00001	1.80	1.4	0.25	scheduled_fixed	6.8–15	8.8 ± 0.6	1.78	5.1

Within-wave load control is useful, but not universally superior to a fixed load. The best overall L1 row uses within-wave control. The best 2.5 m row uses a global fixed load. In these simulations, load tuning acts as both a power-matching tool and an electromagnetic damping adjustment. For the headline L1 rows in Table 9.7, per-coil RMS current remains below 0.17 A.

Strict 1.8 m envelope

The targeted strict-envelope confirmation evaluated LR00001 and LR00008 over the full ten representative states with five seeds per state. Only fixed-load policies were carried into this final strict-package comparison.

The strongest strict-envelope row is LR00001 ($F_c \times 0.25$) with scheduled fixed load and $m_{\text{EM}} = 1.4$. It has $N = 12$, $m = 40$ kg, $f_n = 0.45$ Hz, $k_{\text{eff}} \approx 320$ N/m, and a 1.8 m package. It produced

$$\bar{P}_L = 8.8 \pm 0.6 \text{ W}, \quad L_{\text{rod}} = 1.78 \text{ m}, \quad a_{\text{peak}} \approx 5.1 \text{ g}. \quad (9.10)$$

The peak stop force was approximately 1900 N, the maximum stop-contact fraction was approximately 2.7%, the peak per-coil RMS current was approximately 0.14 A, and no bottom-out occurred. The scheduled loads span $6.8\ \Omega$ to $15\ \Omega$. The corresponding global fixed-load row produced nearly the same result, $\bar{P}_L = 8.76 \pm 0.59\ \text{W}$, with a $10\ \Omega$ load. Sea-state scheduling therefore gives no practically meaningful gain for this strict-envelope candidate, although the wider L1 results still indicate that load adaptation can be useful for other operating points.

At lower electromagnetic multipliers, LR00001 reaches approximately 8.7 W at $m_{\text{EM}} = 1.2$ and 8.2 W at $m_{\text{EM}} = 1.0$, but both rows fail the configured acceleration and bottom-out checks in at least one higher- H_s state. LR00008 did not pass all configured checks in the strict-package confirmation. Its strongest raw rows were near 8.0 W to 8.6 W, but they failed at least one configured check, mainly through bottom-out and package exceedance.

The strict 1.8 m result improves from approximately 3.8 W in O2 and 2.4 W in unoptimised B2 to approximately 8.8 W after load retuning. It remains below the 10 W target on the full ten-state weighting. Closing the gap would require higher electromagnetic coupling, a larger package, a different stroke-management law, or full ten-state evaluation of within-wave policies at the strict package length.




Figure 9.8: Load-tuning recovery from B2 to L1. The LR00001 bar uses the confirmed ten-state strict-envelope result.




Figure 9.9: L1 feasibility landscape. Ocean-weighted external-load power is plotted against worst-case required rod length. Strict-1.8 m confirmation rows are overlaid.




Figure 9.10: Per-sea-state weighted contribution $P_{\text{mean}} w$ for three L1 reference rows: (a) LR00241, (b) LR00169, and (c) strict-envelope LR00001. Operational states are SS01–SS06 and stress states are SS07–SS10.

9.7 Summary of Scaling Results

The scaling analysis does not identify a final 10 W design inside the strict 1.8 m rod-length envelope. It does, however, identify the design conditions that bring the selected architecture close to that target. The main requirements are a large moving mass, improved electromagnetic coupling, reduced Coulomb friction (25–50% of the validated prototype value), retuned electrical loading, and a stroke-management mechanism that limits travel without introducing excessive acceleration or repeated bottom-out.

The free-carriage ocean screens indicate that stroke length is the first major constraint. In O1, the highest raw ocean-weighted power was approximately 15.2 W, but it required $L_{\text{rod}} = 9.34$ m. Even the shortest nonlinear O1 candidate required 5.48 m. When O2 used the 1.8 m package envelope as a primary filter, the best strict free-carriage result fell to approximately 3.8 W. A relaxed-envelope O2 candidate reached approximately 7.5 W at $L_{\text{rod}} = 2.56$ m, but applying the strict 1.8 m cutout reduced its usable power to approximately 3.9 W. The high-power free-carriage cases therefore depend on stroke lengths that are unavailable in the strict package.

Target natural frequency is also constrained by packaging. The strongest low-frequency ocean-screened candidates used low effective stiffness and produced high raw power, but required several metres of relative travel. Within the strict package, the preferred target natural frequency shifts upward. The leading ocean-screened candidates use $f_n \approx 0.45$ Hz to 0.50 Hz, rather than the lower frequencies associated with the dominant ocean-wave periods. This suggests that direct tuning to the main low-frequency wave component is not necessarily beneficial for a compact internal harvester, because the required stroke can outweigh the response gain.

The stroke-management studies reduce the stroke problem but do not fully resolve it. The B1 diagnostic indicates that nonlinear restoring or stop forces can reduce package length while preserving several watts of power. Its absolute values are not directly comparable to the wave-follower ocean screens because B1 uses a cylinder-heave RAO and a different sea-state partition. The integrated B2 bumper sweep is more directly relevant to the final ocean-screened cases. In B2, no final candidate passed all configured controlled-stroke checks. The strongest B2 candidate produced approximately 4.2 W at $L_{\text{rod}} = 2.51$ m, while the best strict-package B2 candidate produced approximately 2.4 W at $L_{\text{rod}} = 1.80$ m. Both were limited mainly by peak acceleration and bottom-out. The tested bumper formulation was therefore insufficient without further operating-point and load retuning.

Electrical load tuning gives the largest improvement in the final scaling sequence. For the strict 1.8 m envelope, the best five-seed result improves from approximately 3.8 W in O2 and 2.4 W in untuned B2 to

$$\bar{P}_L = 8.8 \pm 0.6 \text{ W}$$

after L1 retuning, using LR00001 with $m_{\text{EM}} = 1.4$ and $F_c \times 0.25$. This row passes the configured checks, with $L_{\text{rod}} = 1.78$ m, peak relative carriage acceleration of approximately 5.1 g, peak stop force of approximately 1900 N, maximum stop-contact fraction of approximately 2.7%, and no bottom-out. For this strict-envelope candidate, sea-state scheduling selects a similar load range to the best global fixed-load case, since a single 10 Ω load gives nearly the same ocean-weighted mean power.

The strict-envelope result does not reject load adaptation. Across L1, the electrical load affects both delivered power and electromagnetic damping. The best overall L1 candidate uses a within-wave load-control policy to retain 10 W-class power while satisfying the configured checks, whereas comparable static-load rows produced similar raw power but

failed through bottom-out or excessive acceleration. Future hardware should therefore allow controlled load switching between sea states, with faster within-wave switching as a possible extension.

The 10 W-class results occur outside the strict package or under idealised forcing. The S2 sinusoidal long rerun reaches approximately 10.2 W at $L_{\text{rod}} = 1.76$ m, but only for an imposed 4 s sinusoid with $m_{\text{EM}} = 1.5$. Under irregular wave-follower forcing, LR00241 reaches approximately 10.4 W at $L_{\text{rod}} = 2.69$ m with $m_{\text{EM}} = 1.4$ and $F_c \times 0.25$, while LR00197 reaches approximately 10.0 W at $L_{\text{rod}} = 2.45$ m with the same electromagnetic multiplier and $F_c \times 0.50$. At $m_{\text{EM}} = 1.0$ with $F_c \times 0.25$, LR00169 reaches approximately 9.8 W at the 2.5 m package class. These results support the plausibility of 10 W-class output in the harvester-side model, but not within the strict 1.8 m envelope.

The combined outcome of the scaling sequence is summarised in Figure 9.11. The figure compares the predicted ocean-weighted external-load power with the worst-case required rod length for the unique candidates retained from the screening stages. It therefore shows both parts of the design target at once: whether a candidate reaches the 10 W power class, and whether it remains inside the strict 1.8 m rod-length envelope.




Figure 9.11: Combined feasibility map across all scaling stages. Ocean-weighted external-load power is plotted against worst-case required rod length. The shaded target region is $L_{\text{rod}} \leq 1.8$ m and $\bar{P}_L \geq 10$ W. No candidate lies inside the strict target region.

Figure 9.12 complements the combined feasibility map by showing the design-variable structure behind the same 251 unique plotted candidates. Each line represents one candidate and links the screening stage, main design settings, stroke-management parameters, and binned performance outcomes. This makes it possible to compare the parameter combinations that survived the successive screening stages and to identify the trends associated with the strongest candidates.




Figure 9.12: Explored design space for the 251 unique candidates included in the combined Pareto summary. Each line represents one candidate. The axes show the screening stage, main design parameters, stroke-management settings, and binned outcomes for power and required rod length. Line colour indicates predicted mean external-load power.

The strongest candidates cluster at the upper end of the tested moving-mass range. In the final relevant screens, the best rows use $m = 40$ kg, which was the upper mass limit in the main B2/L1 grid. The sweep therefore does not identify the mass level at which further increases become impractical or give diminishing benefit. It indicates that the leading candidates require a heavy internal moving body, which is a major system-level requirement for a man-portable buoy. This mass must be assessed against hull volume, payload, ballast, handling, deployment, and survivability constraints.

The preferred natural-frequency range shifts between the idealised sinusoidal screens and the irregular ocean screens. The 4 s sinusoidal optimum uses $f_n = 0.25$ Hz, while the strict free-carriage O2 candidate uses $f_n = 0.50$ Hz, corresponding to $k_{\text{eff}} \approx 395$ N/m for $m = 40$ kg. The controlled-stroke L1 candidates mainly use $f_n = 0.45$ Hz, corresponding to $k_{\text{eff}} \approx 320$ N/m. The higher ocean-screened frequencies reduce stroke under broadband irregular forcing. A future restoring mechanism would therefore need to provide an effective carriage stiffness on the order of 300 N/m to 400 N/m for the leading ocean-screened designs, before nonlinear stop or hardening terms are added.

Table 9.8: Best 1.8 m envelope results across the campaign, all on the full ten-state weighting.

Stage	ID	EM	F_c	$\bar{P}_L$	Screening status	Limiting issue or note
		$\times$	$\times$	[W]		
O2 stroke-feasible	OC0134	1.5	0.25	3.8	passes stroke check	free carriage, no stroke management
B2 integrated bumpers	LR00001	1.0	0.25	2.4	does not pass configured checks	peak acceleration and bottom-out
L1 confirmation, $m_{EM} = 1.0$	LR00001	1.0	0.25	8.2 ± 0.4	does not pass configured checks	peak acceleration and bottom-out
L1 confirmation, $m_{EM} = 1.2$	LR00001	1.2	0.25	8.7 ± 0.5	does not pass configured checks	peak acceleration and bottom-out
L1 confirmation, $m_{EM} = 1.4$	LR00001	1.4	0.25	8.8 ± 0.6	passes configured checks	scheduled fixed-load policy

The active constraints are mechanical. Free-carriage ocean results are limited by required rod length. With stroke management active, peak relative carriage acceleration and bottom-out become limiting. Coil current is not limiting in these screens: the largest per-coil RMS current in the ocean-final candidates is approximately 0.17 A.

The stress-state failures suggest that future designs should include a deliberate survival or protection mode. The present bumper model limits travel through passive end-stop forces only, and the failed B2 and L1 rows indicate that bottom-out and peak acceleration can become active constraints in higher- H_s states. A practical buoy design may therefore need a locking, braking, or high-damping mode that protects the carriage during severe seas, even if that mode sacrifices energy capture during those intervals.

The main limitations remain important. The wave-follower input is an upper-bound base-motion model and is not coupled to buoy heave dynamics. The electromagnetic multipliers used by the best L1 rows are design requirements on $d\lambda/dq$, not measured hardware values. The friction multipliers are required reductions in mechanical loss for a future guide and bearing design, not tested hardware values. The stroke-management terms are abstract force laws rather than realised mechanical components. A physical restoring and end-stop mechanism, a justified acceleration limit, adaptive electrical loading, and a buoy-coupled motion model are required before the architecture can be treated as a deployable 10 W design.

Chapter 10

Conclusions and Future Work

10.1 Overview

This thesis investigated a sealed internal wave energy harvester for a self-powered monitoring buoy in Icelandic coastal waters. The selected architecture uses linear relative motion between a translating coil carriage and a segmented magnetic rod. The long-term target was an average electrical load power of approximately 10 W for a man-portable buoy, while the present work focused on concept selection, electromagnetic modelling, prototype construction, experimental characterisation, coupled electromechanical validation, and harvester-side scaling.

The work did not aim to deliver a deployed buoy. It assessed whether the selected internal linear electromagnetic architecture is credible enough for further development and quantified the main design changes required to move from the bench-scale prototype toward the long-term target.

10.2 Summary of the Work

The work began with a concept-selection process for sealed internal buoy harvesters. A linear sliding electromagnetic architecture was selected because it combines mechanical simplicity, direct modelling tractability, low prototype cost, and a clear scaling path through active length, moving mass, electromagnetic coupling, and electrical load tuning.

A segmented LinMot magnetic rod was then reverse-engineered and modelled. Surface magnetic-field measurements were used to infer the internal magnet-spacer pattern and calibrate finite-element field models in ANSYS Maxwell and MATLAB PDE Toolbox. The field reconstruction supported coil-level optimisation, which identified an optimum coil outer diameter close to 48.7 mm for the selected winding constraints. The physical prototype used three 800-turn coils mounted on a translating carriage, with the middle coil reversed in polarity so that the three coil voltages add constructively.

A bench-scale prototype, embedded measurement system, custom coil winder, and rope-driven wave simulator were developed. Free-decay experiments, load sweeps, frequency-response tests, and amplitude-response tests were then used to identify the main mechanical and electromechanical parameters of the prototype. The finite-amplitude natural frequency was approximately 0.83 Hz, the adopted voltage-speed constant was $K_{\text{emf}} \approx 12.7 \text{ V s/m}$, and the electromagnetic damping increased approximately with $1/R_{\text{tot}}$ for fixed K_{emf} , where R_{tot} includes the coil and external circuit resistance. The open-circuit mechanical loss was better represented by a Coulomb-plus-viscous model than by a single viscous coefficient.

The cycle-resolved loss partition, with $c_{\text{visc}} = 3.40 \text{ N s/m}$ and $F_c = 0.95 \text{ N}$, was then used as the fixed mechanical-loss input for the coupled validation.

When measured carriage motion was replayed through the PDE-derived electrical model, using the fixed global phase correction and training-derived electromagnetic scale factor, the predicted run-averaged bypass-path power agreed closely with measurements, with an all-run median power ratio of 1.019. In the full coupled validation, where carriage motion was predicted from measured base excitation, the median power ratio was 0.941, the median carriage-amplitude ratio was 0.785, and the median carriage-motion agreement was $q_{R^2} = 0.905$. These results indicate that the electromagnetic conversion model is well supported when motion is known, while the largest remaining uncertainty in the prototype validation is the predicted mechanical response amplitude.

Finally, applying a wave-follower kinematic assumption to model upper-bound base excitation, the validated electromagnetic model and single-degree-of-freedom harvester framework were used for scaling analysis. The scaling study showed that the selected architecture has a conditional path toward 10 W, but not inside the strict 1.8 m rod-length envelope under the final ten-state ocean-weighted screen. The best strict-envelope result reached approximately $8.8 \pm 0.6 \text{ W}$ assuming a 40 kg moving mass, $m_{\text{EM}} = 1.4$, reduced Coulomb friction, active stroke management, and scheduled fixed electrical loading. The strongest feasible ocean-weighted result reached approximately $10.4 \pm 0.6 \text{ W}$, but required a rod length of 2.69 m alongside the improved electromagnetic coupling.

10.3 Main Conclusions

The main conclusions of the thesis are as follows:

1. **A sealed internal linear electromagnetic architecture is a credible candidate for further development.** Its main advantages are simple internal packaging, direct electromagnetic conversion, modular active length, and tractable one-dimensional modelling. The bench-scale prototype demonstrated watt-level output under controlled resonance-region excitation, with the best tested load (9.8Ω) producing a mean bypass-path power of approximately 4.78 W.
2. **The 10 W target is conditionally plausible, but was not achieved inside the strict 1.8 m package under the modelled conditions.** The best strict-envelope candidate predicted approximately $8.8 \pm 0.6 \text{ W}$. Candidates exceeding 10 W required a larger package of approximately 2.45 m to 2.69 m, a 40 kg moving mass, improved electromagnetic coupling, and active stroke management.
3. **Stroke length, not generation capacity, is the binding design constraint.** Free-carriage numerical screens yielded candidates predicting over 15 W when allowed several metres of relative travel, while the same architecture compressed into the strict 1.8 m envelope yielded only 3.8 W without further intervention. The 10 W target is therefore primarily a packaging problem, not a problem of insufficient electromagnetic conversion. Coil current was not limiting in any leading candidate (per-coil RMS current remained below 0.17 A); rather, the active constraints were required stroke, peak relative carriage acceleration, end-stop impact dynamics, and the practical implementation of a 40 kg moving mass.
4. **Directly tuning to the dominant wave-frequency range is not beneficial under a strict package constraint.** The idealised 4 s sinusoidal optimum sat at $f_n =$

0.25 Hz, which lies at the higher-frequency end of typical ocean-wave spectra. Under irregular wave-follower forcing within the 1.8 m envelope, however, the strongest candidates use $f_n \approx 0.45$ Hz to 0.50 Hz, corresponding to effective stiffnesses of approximately 300 N/m to 400 N/m for a 40 kg moving mass. Lower target natural frequencies increase relative motion but exceed the available stroke. The preferred operating point for a compact internal harvester therefore lies above, not at, the dominant wave frequency.

5. **Electrical load tuning is the single largest recoverable design lever.** Within the strict 1.8 m envelope, load retuning alone raised the ocean-weighted result from approximately 2.4 W (untuned controlled-stroke screen) to approximately 8.8 ± 0.6 W using the same geometry. Load resistance acts simultaneously as a power-matching parameter and as an electromagnetic damping adjustment, so the optimum is a coupled electromechanical result rather than a simple resistance-matching result. Within-wave load control retained 10 W-class power at larger packages while satisfying acceleration and end-stop impact constraints that comparable fixed-load configurations failed.
6. **The prototype mechanical loss is not purely viscous, and the loss decomposition matters quantitatively for forward prediction.** Multi-amplitude open-circuit decay tests showed that a Coulomb-plus-viscous representation matched behaviour better than a single viscous coefficient. The cycle-resolved decomposition ($c_{\text{visc}} = 3.40$ N s/m, $F_c = 0.95$ N) raised the median coupled-validation power ratio from 0.523 to 0.941 compared with an envelope-fit alternative carrying the same total free-decay loss, showing that how the loss is partitioned, not only its magnitude, controls forward-prediction accuracy in the off-resonance regime.
7. **The standalone electromagnetic model exhibits higher predictive accuracy than the fully coupled dynamic model.** When measured carriage motion was prescribed, the PDE-derived electrical model predicted mean bypass-path power with an all-run mean ratio of 1.013 and a zero-intercept slope of 0.992. The full coupled model reproduced the dominant carriage waveform and phase (median $q_{R^2} = 0.905$, median fundamental phase delay -2.1°), but the response amplitude remained regime-dependent. The validation therefore supports the electromagnetic conversion model strongly and the coupled mechanical response more conditionally. Because the forward scaling analysis abstracts the restoring mechanism as an idealised stiffness search, the electrical conversion step is well-supported, whereas the scaled carriage motion remains an idealised projection.
8. **Harmonic excitation actively drives off-resonance power generation.** When subjected to long-period base motion, higher harmonics of the excitation can align with the prototype's natural frequency and become amplified. This produces measurable carriage motion and bypass-path power even when the fundamental forcing frequency remains well below resonance. This mechanism is highly relevant for operation in stochastic seas; because real ocean waves provide broadband excitation, higher-frequency spectral components can actively excite the intentionally detuned harvester near its natural frequency.

10.4 Answers to the Research Questions

RQ1: What is the most suitable and cost-efficient internal mechanical architecture for harvesting kinetic energy within a sealed, volume-constrained buoy enclosure?

The linear sliding electromagnetic architecture was the most suitable and cost-efficient choice for the scope of this thesis. It scored highest in the concept-selection process because it combines low mechanical complexity, low prototype cost, sealed internal implementation, strong modelling tractability, and a clear scaling path through active length, moving mass, electromagnetic coupling, and electrical loading.

RQ2: How can the selected architecture be represented by a coupled electromagnetic and mechanical model, and what predictions does the model make about its behaviour?

The selected architecture can be represented as a one-degree-of-freedom base-excited oscillator coupled to a rod-coil electrical model through an electromagnetic reaction force. The model predicts coupled electromechanical behaviour rather than a purely mechanical oscillator response: response and power increase near resonance, the nonlinear rope restoring force is predicted to produce amplitude-dependent resonance, and electrical loading changes the mechanical response through electromagnetic damping, approximately following $c_{em} = K_{emf}^2/R_{tot}$. The optimum external load is therefore predicted to be a coupled dynamic outcome rather than a direct match to the coil resistance.

RQ3: How well does the model reproduce the measured electrical and mechanical behaviour of the physical prototype under controlled laboratory excitation?

The experimental results support most of the model's qualitative predictions, but the fixed-parameter coupled model has mixed quantitative accuracy. The measured electromagnetic damping followed the expected $1/R_{tot}$ trend with $R^2 \approx 0.997$, the load sweep produced maximum bypass-path power at an intermediate external load, and the frequency-response experiments located the resonance peak within 1% of the decay-identified natural frequency. The amplitude-dependent resonance prediction was not supported over the reliable tested range, where the measured natural frequency showed no statistically significant amplitude dependence.

Quantitatively, the electromagnetic replay model predicted mean bypass-path power well when measured carriage motion was prescribed, with an all-run mean power ratio of 1.013 and a zero-intercept slope of 0.992. The full coupled model reproduced waveform and phase reasonably well, but response amplitude remained the main source of uncertainty.

RQ4: Based on the validated model, what design changes are required for the concept to move toward the target average power level of approximately 10 W?

The scaling analysis indicates that approaching the 10 W target requires several simultaneous design changes: increased moving mass, improved electromagnetic coupling, reduced mechanical loss, controlled electrical loading, and a physical restoring and stroke-protection system. The strongest candidates use approximately 40 kg moving mass, which is substantial for a man-portable buoy. They also require an effective electromagnetic coupling increase of about 40% above the prototype baseline and, under irregular wave-follower forcing, package lengths of approximately 2.45 m to 2.69 m for 10 W-class operation. Within the strict 1.8 m envelope, the target is not reached, although load retuning increased the best controlled-stroke result to approximately (8.8 ± 0.6) W.

Future hardware should support sea-state-dependent effective-load scheduling, with faster within-wave control as a possible extension. The scaled design also requires a physically implemented stroke-management system. Compliant bumpers, end stops, controllable braking, high-damping operation, or a locking mode may be needed to limit bottom-out and peak carriage acceleration in severe sea states. The required base motion also needs to be assessed against a buoy-coupled hydrodynamic model, since the present scaling treats the wave elevation as an upper-bound input. The remaining bottlenecks are mainly mechanical integration: package length, stroke control, acceleration limits, and protection during high-energy sea states.

10.5 Limitations

The conclusions above should be interpreted within the following limitations.

Modelling scope.

- **Harvester-side model only:** The dynamic model represents the internal harvester as a one-degree-of-freedom oscillator driven by prescribed base motion. It does not include hull hydrodynamics, six-degree-of-freedom buoy response, mooring dynamics, or the reaction force from the harvester back on the hull. The wave-follower assumption used in the scaling analysis assigns the full wave elevation to the harvester base and is therefore an optimistic upper-bound excitation.
- **Abstract restoring and stroke-management laws:** The scaled cases prescribe an effective linear stiffness with optional cubic hardening and a polynomial end-stop, rather than extrapolating the prototype elastic-rope mechanism. These force laws represent design requirements for a future restoring and stroke-protection system; their validation uncertainty does not transfer directly from the prototype rope system. The 6 g peak acceleration threshold used in the controlled-stroke screens is a provisional screening limit, not a derived survivability requirement.
- **Simplified electrical output:** Reported power is dissipated in resistive loads. The model does not include rectification, energy storage, conversion efficiency, or sensor-load scheduling losses, all of which would reduce the net stored power in a deployed buoy.

Validation scope.

- **Calibrated rather than predicted dynamic stiffness:** A fixed rope-force scale factor of 1.634 was applied so that the model reproduced the experimentally identified dynamic natural frequency rather than the static rope law. This calibration choice resolves the dominant timescale but does not guarantee correct stiffness across all loading conditions and contributes to the regime-dependent response-amplitude bias seen in the coupled validation.
- **Carriage-motion estimation:** The carriage translation signal is reconstructed from a Hall-counted position anchor and magnetometer phase interpolation, and is subject to drift over long runs. This limits how closely the coupled validation can match measured carriage motion at the time-series level.
- **Wave-simulator tracking:** The wave simulator did not act as an ideal displacement source at the lowest commanded amplitudes, achieving 45% of the commanded amplitude at $A_{\text{cmd}} = 0.2$ m. Analyses use the achieved rather than the commanded amplitude, so this limitation appears in the measured excitation rather than as an analysis error.

Forward-scaling scope.

- **Unvalidated design targets:** The leading scaled candidates require an effective electromagnetic coupling multiplier of $m_{\text{EM}} = 1.4$ and Coulomb-friction multipliers below unity. These represent design requirements for future hardware iterations, not measured improvements.
- **Representative sea states:** The ten-state ocean weighting captures the distribution of (H_s, T_z) at the Surtseyjardufl record but does not preserve storm sequencing, calm periods, wave-group statistics, or site-to-site coastal variation. The weighted-power metric estimates expected output over the representative distribution rather than continuous annual operation.
- **Design-space coverage and resolution:** The scaling analysis explored a defined grid of moving mass, coil count, target natural frequency, electrical load, electromagnetic multiplier, friction multiplier, package length, and stroke-management settings across multiple stages. Computing time was the practical limit; some sweeps required more than 15 hours, so the parameter grid was kept coarser than would be ideal and not every combination was retested at every later stage. Better designs may therefore exist within the explored region at finer resolution, and parameter combinations outside the explored region were not tested. In particular, the leading candidates' 40 kg moving-mass requirement could plausibly be relaxed at higher electromagnetic coupling than $m_{\text{EM}} = 1.4$. Whether such coupling is achievable in practice is itself a question for future electromagnetic redesign, which Chapter 4 identifies as a likely route through changes such as laminated soft-magnetic flux concentrators, reduced air gaps, or revised magnetic-circuit geometry.

10.6 Recommended Future Work

The next development steps should focus on the parts of the system that currently limit confidence in the 10 W path.

Buoy-coupled hydrodynamic model

The most important next step is to replace the wave-follower base-motion assumption with a buoy-coupled model. This model should include hull heave response, added mass, radiation damping, excitation force, hydrostatic restoring, mooring influence, and the reaction force from the internal harvester. The output should be the actual internal base motion available to the harvester for candidate buoy geometries. This step is necessary before the ocean-weighted 10 W claim can be treated as a buoy-level prediction.

Electromagnetic redesign and hardware validation

The best scaled candidates require electromagnetic coupling above the measured prototype baseline. A dedicated electromagnetic redesign should therefore be performed. Candidate improvements include reducing the coil-magnet air gap, increasing active magnetic loading, improving flux guidance, adding laminated or segmented soft-magnetic cores around the coils, redesigning the winding geometry, increasing conductor fill, reducing resistance, and evaluating alternative magnet arrays. The target should be to demonstrate the required m_{EM} experimentally, not only in simulation.

Physical stroke-management mechanism

The scaling analysis shows that stroke management is essential. Future work should design and test a physical end-of-travel system that limits stroke without producing excessive acceleration or impact force. Possible approaches include progressive springs, magnetic end stops, elastomeric buffers, hydraulic or eddy-current damping, controllable dampers, mechanisms that redirect excess stroke into additional electrical generation, or an active locking mode for severe sea states. A survival or protection mode may be useful even if energy capture is reduced during severe states, because peak relative carriage acceleration and bottom-out are the main failure modes in the controlled-stroke and load-tuning results.

Mechanical-loss reduction

Mechanical-loss reduction is another major design lever. A future prototype should replace the present manual-fit guide arrangement with an engineered linear bearing system. Rod tolerances, alignment, sealing-compatible lubrication, guide preload, moving-wire routing, and environmental contamination should be treated as design variables. The next prototype should include direct measurement of guide friction under representative loads and environmental conditions.

Integrated load-control electronics

The L1 results show that load selection has a strong influence on power and feasibility. Future work should implement an electrical interface capable of controlled effective load variation. This could begin with fixed-load switching or sea-state scheduling and then progress

to a real-time power-electronics stage. The controller should be evaluated by its net stored energy after conversion losses, not only by instantaneous external-load power.

Scaled prototype

A second-generation prototype should combine the strongest lessons from the thesis: increased moving mass, improved electromagnetic coupling, reduced mechanical loss, a controlled stroke-management mechanism, and a realistic power-electronics interface. The prototype should be sized to test the critical 1.8 m to 2.7 m package range identified in Chapter 9. It should be tested first on a more capable controlled-motion platform than the bench-scale wave simulator used in this thesis, and then in sheltered coastal water before open-ocean deployment.

Long-duration and environmental testing

For a deployable monitoring buoy, the sealed internal mechanism must survive repeated cycling, temperature variation, vibration, shock, humidity, leakage risk, corrosion, fouling, and long-duration operation. Future testing should include endurance cycling of the guide system, ropes or springs, stops, bearings, electrical connections, and moving-wire interfaces. These tests are necessary to convert the present proof-of-concept architecture into a reliable marine subsystem.

10.7 Final Remarks

This thesis demonstrates that a sealed internal linear electromagnetic harvester can produce useful watt-level power. The harvester can be modelled with experimentally validated electromagnetic and coupled dynamic models, and can be scaled systematically to identify the design requirements for a practical buoy-scale system.

The bench-scale prototype showed watt-level output under controlled excitation near resonance. When measured motion was prescribed, the electromagnetic model reproduced measured power closely, and the coupled model reproduced the main mechanical response trends well enough to support forward screening. The scaling analysis then showed that the path to the 10 W class is technically plausible but demanding. It requires a heavy internal moving mass, improved electromagnetic coupling, reduced mechanical loss, careful electrical load tuning, and a practical stroke-management system.

Within the strict 1.8 m package, the best final screen predicts approximately 8.8 ± 0.6 W. With a larger package of approximately 2.45 m to 2.69 m, the model prediction reaches the 10 W class. The thesis therefore does not present a fully integrated self-powered buoy system, but it does narrow the problem from a broad concept to a better-defined set of engineering requirements.

The selected architecture is worth further development. The next phase should move from harvester-side screening to buoy-integrated design, combining hydrodynamic modelling of the buoy response with the internal harvester model developed in this thesis. The decisive questions are how the required buoy motion, moving mass, package length, electromagnetic coupling, stroke control, and electrical interface can be integrated into a practical man-portable buoy. In that sense, the main outcome of this work is a validated foundation for the next design stage: turning the sealed internal harvester concept into a practical wave-powered buoy.

Data, Code, and Supporting Materials

The data, code, and supporting files for this thesis are available in the project repository:

<https://github.com/HafthorArni/wave-energy-harvester-thesis>

The repository contains the experimental datasets, MATLAB and Python analysis scripts, generated figures and numerical outputs, CAD material, and selected hardware and firmware snapshots used during the project. It also includes supplementary reports for the segmented magnetic rod, the bench-scale harvester prototype, and the rope-driven wave-simulator platform. These reports provide additional implementation detail for the physical systems that are described more briefly in the thesis.

The repository is intended to make the main experimental, modelling, and scaling workflows traceable and reproducible from the supporting files.

Bibliography

- [1] Y. Li et al., “Study of an electromagnetic ocean wave energy harvester driven by an efficient swing body toward the self-powered ocean buoy application,” *IEEE Access*, vol. 7, pp. 129 758–129 769, 2019. DOI: 10.1109/ACCESS.2019.2937587.
- [2] A. Lo Schiavo, F. Nicora, and C. Boragno, “Modeling a wave energy harvester for supplying data buoys,” *Applied Sciences*, vol. 14, no. 16, p. 7019, 2024. DOI: 10.3390/app14167019.
- [3] H. Jung et al., “Self-powered ocean buoy using a disk-type triboelectric nanogenerator with a mechanical frequency regulator,” *Nano Energy*, vol. 121, p. 109 216, 2024. DOI: 10.1016/j.nanoen.2023.109216.
- [4] T. Zhao et al., “A pendulum-based nanogenerator for high-entropy wave energy harvesting,” *Nature Communications*, vol. 16, p. 5480, 2025. DOI: 10.1038/s41467-025-60443-8.
- [5] J. T. Cheung and E. F. I. Childress, “Ocean wave energy harvesting devices,” Teledyne Scientific & Imaging, LLC, Tech. Rep. Final Report RC71273.FR, 2007.
- [6] M.-C. Chiu, M. Karkoub, and M.-G. Her, “Energy harvesting devices for subsea sensors,” *Renewable Energy*, vol. 101, pp. 1334–1347, 2017. DOI: 10.1016/j.renene.2016.10.018.
- [7] J. Cortés et al., “Open sea lab: An integrated coastal ocean observatory powered by wave energy,” *Journal of Marine Science and Engineering*, vol. 10, no. 9, p. 1249, 2022. DOI: 10.3390/jmse10091249.
- [8] M. H. Maheen and Y. Yang, “Wave energy converters with rigid hull encapsulation: A review,” *Sustainable Energy Technologies and Assessments*, vol. 57, p. 103 273, 2023. DOI: 10.1016/j.seta.2023.103273.
- [9] J. Qin, Z. Zhang, X. Song, S. Huang, Y. Liu, and G. Xue, “Design and performance evaluation of an enclosed inertial wave energy converter with a nonlinear stiffness mechanism,” *Journal of Marine Science and Engineering*, vol. 12, no. 1, p. 191, 2024. DOI: 10.3390/jmse12010191.
- [10] R. Ahamed, K. McKee, and I. Howard, “A review of the linear generator type of wave energy converters’ power take-off systems,” *Sustainability*, vol. 14, no. 16, p. 9936, 2022. DOI: 10.3390/su14169936.
- [11] G. Bracco, E. Giorcelli, G. Mattiazzo, F. Marignetti, S. Carbone, and C. Attaiatese, “Design and experiments of linear tubular generators for the inertial sea wave energy converter,” in *Proc. IEEE Energy Conversion Congress and Exposition (ECCE)*, 2011. DOI: 10.1109/ECCE.2011.6064294.

- [12] L. Cappelli et al., “Linear tubular permanent-magnet generators for the inertial sea wave energy converter,” *IEEE Transactions on Industry Applications*, vol. 50, no. 3, pp. 1817–1828, 2014. DOI: 10.1109/TIA.2013.2291939.
- [13] T. D. Dang, T. C. Do, and K. K. Ahn, “Design and experiment of a direct-drive wave energy converter using a variable stiffness mechanism,” *Ocean Engineering*, vol. 266, p. 112 841, 2022. DOI: 10.1016/j.oceaneng.2022.112841.
- [14] R. B. Ummaneni, J. E. Brennvall, and R. Nilssen, “Convert low frequency energy from wave power plant to high frequency energy in linear electrical generator with gas springs,” in *Proc. 2008 Joint International Conference on Power System Technology and IEEE Power India Conference*, 2008. DOI: 10.1109/ICPST.2008.4745370.
- [15] H. Chen, N. Liu, and X. Liu, “Wave energy linear generator system including a newly designed wave roller mechanical interface,” *Actuators*, vol. 14, no. 3, p. 147, 2025. DOI: 10.3390/act14030147.
- [16] T. Wang and H. Wang, “Multistable pendulum wave energy harvesting under multi-directional irregular excitations,” *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 4, pp. 3175–3183, 2024. DOI: 10.1109/TMECH.2024.3400989.
- [17] S. Crowley, R. Porter, D. J. Taunton, and P. A. Wilson, “Modelling of the witt wave energy converter,” *Renewable Energy*, vol. 115, pp. 159–174, 2018. DOI: 10.1016/j.renene.2017.09.031.
- [18] S.-M. Jang, J.-Y. Choi, S.-H. Lee, H.-W. Cho, and W.-B. Jang, “Analysis and experimental verification of moving-magnet linear actuator with cylindrical halbach array,” *IEEE Transactions on Magnetics*, vol. 40, no. 4, pp. 2068–2070, 2004. DOI: 10.1109/TMAG.2004.832157.
- [19] O. Farrok, M. R. Islam, M. R. I. Sheikh, Y. Guo, J. Zhu, and G. Lei, “Oceanic wave energy conversion by a novel permanent magnet linear generator capable of preventing demagnetization,” *IEEE Transactions on Industry Applications*, vol. 54, no. 6, pp. 6005–6014, 2018. DOI: 10.1109/TIA.2018.2863661.
- [20] R. Crozier, H. Bailey, E. Spooner, P. McKeever, and M. A. Mueller, “Analysis, design and testing of a novel direct-drive wave energy converter system,” in *Proc. IET Conference on Renewable Power Generation (RPG 2011)*, Edinburgh, UK, 2011. DOI: 10.1049/cp.2011.0122.
- [21] R. Liu, J. Yu, and W. Hu, “An oscillating water column wave energy converter with a linear generator as power take-off based on magnetic fluid reciprocating seal and air spring,” *Energy Conversion and Management*, vol. 351, p. 121 016, 2026. DOI: 10.1016/j.enconman.2025.121016.
- [22] E. Katsidoniotaki, Y.-H. Yu, and M. Göteman, “Midfidelity model verification for a point-absorbing wave energy converter with linear power take-off,” *International Marine Energy Journal*, vol. 5, no. 1, pp. 67–75, 2022. DOI: 10.36688/imej.5.67–75.
- [23] C. Jin, H. Kang, M. Kim, and F. P. Bakti, “Performance evaluation of surface riding wave energy converter with linear electric generator,” *Ocean Engineering*, vol. 209, p. 108 141, 2020. DOI: 10.1016/j.oceaneng.2020.108141.
- [24] H. F. Guðmundsson, “Búnaður til rafmagnsframleiðslu í bauju,” B.Sc. thesis, Reykjavík University, 2015.
- [25] H. S. Olgeirsson, “Self-sustaining power supply for a low-power sea buoy,” M.Sc. thesis, Reykjavík University, 2021.

- [26] D. S. Stutts, *Equivalent viscous damping*, Supplementary course notes, Missouri University of Science and Technology, Revised 2013, 2009.
- [27] S. S. Rao, *Mechanical Vibrations in SI Units*, 6th ed. Pearson, 2018.
- [28] Copernicus Marine Service, *Arctic ocean–in situ near real time observations*, NetCDF-4 data product, Product ID: INSITU_ARC_PHYBGCWAV_DISCRETE_MYNRT_013_031, 2026. DOI: 10.48670/moi-00031.
- [29] P. H. Oleinik, G. P. Tavares, B. N. Machado, and L. A. Isoldi, “Transformation of water wave spectra into time series of surface elevation,” *Earth*, vol. 2, no. 4, pp. 997–1005, 2021. DOI: 10.3390/earth2040059.
- [30] J. M. J. Journée and W. W. Massie, *Offshore Hydromechanics*, First. Delft, The Netherlands: Delft University of Technology, 2001.