Transactions on Environment, Energy and Earth Sciences
ISSN: 2960-1924, eISSN: 2960-2300 | Volume 7
GTREE 2026

Spectral Co-Design of Ship Rudder Geometry and Heading Autopilot for Energy Efficiency and Manoeuvring Safety

Fuyong Zhou
StormGeo, Shanghai, China
zfyego@ieee.org

Abstract. Rudders determine the manoeuvring safety of a ship and permanently consume propulsive power through profile drag, weather helm and autopilot action. Because rudder geometry and course-keeping controller are designed separately and the energy cost of steering requires long stochastic simulations, their joint trade-off between energy efficiency and safety has remained unquantified. This paper proposes a spectral co-design method that optimises rudder area, span and section thickness together with the parameters of a PID heading autopilot. A baseline-anchored, stall-aware rudder model preserves the captive-model calibration of the MMG standard method while extrapolating to new geometries with section data from a physics-informed airfoil model. A frequency-domain evaluator combines steady-sailing equilibria in wind and wave drift with closed-loop second moments under gusts, slowly varying drift and first-order wave loads, and obtains the mean steering resistance from a second-order expansion of the surge balance. The constrained three-objective problem—rudder-related power, utilisation of the IMO manoeuvring criteria and heading deviation, subject to adverse-weather course keeping, steering-gear rate, cavitation and stability—is solved for the KVLCC2 tanker over a climate derived from 194,475 NOAA buoy records. The evaluator reproduces nonlinear time-domain statistics with $R2 = 0.976–0.996$ at an 1100-fold speed-up. Against the original rudder with a conventionally tuned autopilot, the knee co-design reduces the rudder-related power by 15.0% while improving both safety objectives. In every run, co-design outperforms sequential design and geometry optimisation with a fixed autopilot (hypervolume +2.0% and +2.8%), and nonlinear time-domain simulation over the full climate confirms the saving of the knee design (15.4%). A quantile-based robust formulation raises the Monte Carlo reliability of the Pareto designs from 0.77 to 0.91 under hydrodynamic-coefficient uncertainty.

Keywords: Rudder design; Ship manoeuvrability; Energy efficiency; Multi-objective optimisation; Autopilot co-design; MMG model; Steering losses; KVLCC2.

1. Introduction

The tightening of energy-efficiency regulation for merchant ships has turned every source of avoidable resistance into a design variable. At the same time, reductions in installed power and appendage size must not erode the ability of a ship to turn, to check a yaw motion and to keep course in heavy weather. The International Maritime Organization (IMO) regulates the first two abilities through the Standards for ship manoeuvrability, Resolution MSC.137(76) [1], and the third through the guidelines on minimum propulsion power to maintain manoeuvrability in adverse conditions, MEPC.1/Circ.850/Rev.3 [2]. The research programme that preceded these guidelines showed that energy efficiency and manoeuvring safety must be treated jointly rather than sequentially [3, 4], and recent reviews of manoeuvrability standards point in the same direction [5].

The rudder sits exactly at this intersection. It is the primary steering device of a single-screw merchant ship and therefore determines turning, initial turning and yaw-checking ability [6, 7]. It is also a permanent consumer of propulsive power through three mechanisms: the profile drag of the rudder blade in the accelerated propeller slipstream, the steady “weather helm” needed to balance mean wind and wave-drift moments, and the continuous corrective action of the autopilot in a seaway. Koyama [8] formulated the purpose of an autopilot as the minimisation of the resistance increase caused by steering, and Reid [9] derived steering criteria from the nonlinear surge equation. The magnitude of

Content from this work may be used under the terms of CC BY-NC 4.0 licence (https://creativecommons.org/licenses/by-nc/4.0/).
Published by Warwick Evans Publishing.

these losses depends simultaneously on the rudder geometry, which sets the force per degree of helm and the profile drag, and on the controller, which sets how far and how fast the rudder moves.

In current practice these two parts of the rudder system are designed separately. The rudder area and section are fixed early from empirical guidance and manoeuvring predictions [6, 10]; the autopilot is tuned later for the installed rudder [11]. Four technical bottlenecks prevent a joint, quantitative trade-off between energy efficiency and manoeuvring safety:

Coupling. The closed-loop behaviour depends on the product of controller gain and rudder effectiveness, and the steering-gear rate limit and wave filtering couple both further. Optimising the geometry with a fixed controller, or the controller for a fixed geometry, explores only a slice of the joint design space.

Cost of the energy metric. Steering-induced resistance is a statistical quantity. Estimating it by time-domain simulation requires hours of simulated sailing per sea state and heading, which becomes prohibitive when thousands of candidate designs must be assessed over an operational climate.

Geometry extrapolation. Manoeuvring models calibrated with captive-model tests are commonly extrapolated to new rudder geometries with an empirical linear normal-force law alone [12, 13]. Such a law contains no section drag, which determines the energy cost of a rudder, and no check of the stall and cavitation limits that bound aggressive planforms and thin sections.

Operating conditions and uncertainty. Designs are usually assessed in a single design sea state and with nominal hydrodynamic coefficients, although the coefficients carry appreciable uncertainty [14] and the operational climate is broad.

This paper addresses these bottlenecks with a spectral co-design method for the rudder system of a large tanker. The technical contributions are:

A baseline-anchored, geometry-resolved rudder model that preserves the captive-model calibration of the MMG standard method [13] for the original rudder and extrapolates to new span, area and section thickness through Fujii's aspect-ratio law [12], RANS-derived section normal-force slopes [10] and section drag, stall and pressure distributions from a physics-informed airfoil model [15]. It includes a stall model and a cavitation-inception criterion.

A spectral evaluator of rudder-related power that combines a steady-sailing equilibrium in mean wind and wave drift [16] with closed-loop second moments under wind gusts, slowly varying wave drift [17] and first-order wave loads. The mean steering resistance follows in closed form from a second-order expansion of the surge balance. Against nonlinear time-domain simulation the evaluator reproduces heading and rudder statistics with $R^2 > 0.99$ and the steering resistance with $R^2 = 0.98$, at about three orders of magnitude lower cost.

A constrained three-objective co-design formulation that minimises rudder-related power, the utilisation of the IMO manoeuvring criteria and the heading deviation simultaneously over rudder planform, section and autopilot parameters, subject to adverse-weather course keeping, steering-gear rate, cavitation and stability constraints. The objectives are evaluated over an operational climate derived from 194,475 NOAA buoy sea-state records [18], and a quantile-based robust variant accounts for hydrodynamic-coefficient uncertainty.

A comprehensive experimental assessment: validation against MARIN free-running model tests [19, 20], comparison with a conventional reference design and with sequential design at equal computational budget, four multi-objective evolutionary algorithms and random search, ablation studies, Sobol sensitivity analysis and Monte Carlo robustness verification.

Figure 1 gives an overview of the method. Section 2 reviews related work. Section 3 presents the ship, rudder, environment and climate models. Section 4 derives the spectral evaluator, and Sect. 5 the co-design problem. Section 6 describes the experimental protocol, Sect. 7 reports the results and Sect. 8 discusses their implications. Section 9 concludes the paper.

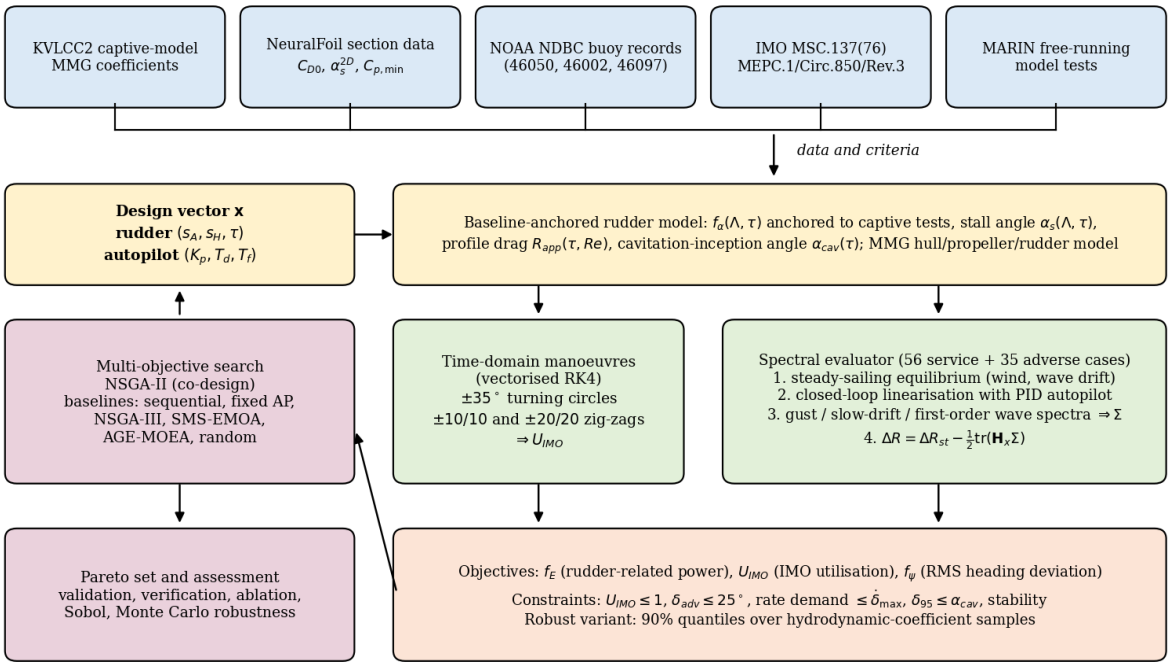


Figure 1. Overview of the spectral co-design method: real data and regulatory criteria feed a baseline-anchored rudder model; each candidate rudder–autopilot design is evaluated by vectorised time-domain manoeuvres and by the spectral evaluator, and a constrained multi-objective search returns the Pareto set, which is assessed by validation, verification, ablation, sensitivity and robustness analyses

2. Related Work

2.1. Rudder Design and Manoeuvring Prediction

System-based manoeuvring prediction separates hull, propeller and rudder forces. The MMG standard method of the Japan Society of Naval Architects and Ocean Engineers formalised the model structure, the captive-model test procedure and the analysis method, and published a complete coefficient set for the KVLCC2 tanker [13]. The KVLCC2 is also the tanker benchmark of the SIMMAN workshops, for which free-running model tests performed at MARIN are available [14, 19]; these data have been used repeatedly to validate captive-test and CFD-based MMG models [10, 20]. In the MMG model the rudder normal force is proportional to a lift-gradient coefficient that is usually estimated from Fujii's formula, which depends only on the aspect ratio [12]. Liu et al. [10] replaced this coefficient by RANS-based regressions for NACA, wedge-tail and IFS profiles and quantified the effect of the profile on the turning and zig-zag indices of the KVLCC2, while assuming that the stall angle lies beyond the range of effective angles. Reviews by Liu and Hekkenberg [7] and by Molland and Turnock [6] summarise how aspect ratio, profile and the propeller slipstream affect rudder lift, drag and stall. These studies analyse given rudders; they do not search the design space and they do not account for the energy consumed by the rudder in service.

2.2. Steering Losses and Course-Keeping Control

Koyama [8] expressed the purpose of an autopilot as the minimisation of the resistance increase caused by steering and derived a quadratic performance index in yaw error and rudder angle; Reid [9] obtained steering criteria from the nonlinear surge equation and designed controllers that minimise added resistance due to steering. Modern guidance-and-control texts present PID heading autopilots with wave filtering and discuss Koyama-type criteria [11]. In these works the rudder geometry is fixed. For manoeuvring in wind and waves, Yasukawa and Sakuno [16] computed steady sailing conditions and course stability with the MMG model, and the SHOPERA project and subsequent

work defined criteria and assessment procedures for manoeuvrability in adverse conditions [3, 4]. Mean wave-drift loads on ships are commonly estimated with the short-wave asymptotic approach of Faltinsen et al. [21] corrected for finite draft [22], and slowly varying drift loads with Newman's approximation [17].

2.3. Multi-objective and Robust Design

Evolutionary multi-objective algorithms such as NSGA-II [23], NSGA-III [24], SMS-EMOA [25] and AGE-MOEA [26] are standard tools for engineering design, and frameworks such as pymoo [27] make them readily available. The hypervolume indicator [28] is the usual measure of front quality. Robust design optimisation accounts for uncertain parameters through expectation, worst-case or quantile measures [29]. Variance-based global sensitivity analysis [30, 31] identifies the variables that drive the objectives. To the author's knowledge, no previous study has optimised the rudder geometry and the course-keeping controller jointly for propulsive energy and IMO manoeuvring safety under a measured operational climate, and none has made such an optimisation tractable through a verified frequency-domain energy evaluator.

3. Modelling

3.1. Ship and Manoeuvring Model

The study ship is the KVLCC2 tanker in full-load condition: $L_{pp} = 320$ m, $B = 58$ m, $d = 20.8$ m, displacement volume $\nabla = 312{,}622$ m^3, $x_G = 11.2$ m, propeller diameter $D_P = 9.86$ m and a semi-balanced horn rudder with span $H_{R0} = 15.8$ m and movable area $A_{R0} = 112.5$ m^2 [20]. The service speed is $U_0 = 15.5$ kn ($Fn = 0.142$). The three-degree-of-freedom MMG equations with origin at midship read [13]

$$\begin{aligned} (m + m_x)\dot{u} - (m + m_y)vr - x_G m r^2 &= X_H + X_R + X_P + X_E, \\ (m + m_y)\dot{v} + (m + m_x)ur + x_G m\dot{r} &= Y_H + Y_R + Y_E, \\ (I_{zG} + x_G^2 m + J_z)\dot{r} + x_G m(\dot{v} + ur) &= N_H + N_R + N_E, \end{aligned} \tag{1}$$

Where the subscripts H, R, P and E denote hull, rudder, propeller and environmental loads. The hull forces are third-order polynomials in the non-dimensional sway velocity $v' = v/U$ and yaw rate $r' = rL_{pp}/U$:

$$\begin{aligned} X'_H &= -R'_0 + X'_{vv}v'^2 + X'_{vr}v'r' + X'_{rr}r'^2 + X'_{vvvv}v'^4, \\ Y'_H &= Y'_v v' + Y'_r r' + Y'_{vvv}v'^3 + Y'_{vvr}v'^2 r' + Y'_{vrr}v'r'^2 + Y'_{rrr}r'^3, \\ N'_H &= N'_v v' + N'_r r' + N'_{vvv}v'^3 + N'_{vvr}v'^2 r' + N'_{vrr}v'r'^2 + N'_{rrr}r'^3. \end{aligned} \tag{2}$$

The propeller thrust is $X_P = (1 - t_P)\rho n^2 D_P^4 K_T(J_P)$ with $J_P = u(1 - w_P)/(nD_P)$, the quadratic open-water polynomial $K_T(J_P)$ and the wake model $(1 - w_P)/(1 - w_{P0}) = 1 + [1 - \exp(-C_1|\beta_P|)](C_2 - 1)$, $\beta_P = \beta - x'_P r'$. The rudder forces are

$$X_R = -(1 - t_R)F_N \sin\delta, \quad Y_R = -(1 + a_H)F_N \cos\delta, \quad N_R = -(x_R + a_H x_H)F_N \cos\delta \tag{3}$$

With the normal force $F_N = 1/2\,\rho A_R U_R^2 C_N(\alpha_R)$, the effective inflow angle $\alpha_R = \delta - \tan^{-1}(v_R/u_R)$, the lateral inflow $v_R = U\gamma_R(\beta - l'_R r')$ and the longitudinal inflow

$$u_R = \varepsilon u(1 - w_P)\sqrt{\eta\left\{1 + \kappa\left(\sqrt{1 + 8K_T/(\pi J_P^2)} - 1\right)\right\}^2 + (1 - \eta)}, \qquad \eta = D_P/H_R \tag{4}$$

All coefficients are the captive-model values of Yasukawa and Yoshimura [13] (Appendix A), with the open-water thrust polynomial of the same propeller [20]. The equations are integrated at full scale with Froude-scaled coefficients, and the propeller revolutions are set at the self-propulsion point of the service speed, as prescribed for the SIMMAN free-running tests. Steering follows a rate-limited actuator with $\dot{\delta}_{\max} = 2.34°/\mathrm{s}$, the Froude-scaled rate of the model tests [10], and $|\delta| \leq 35°$.

3.2. Baseline-Anchored, Geometry-Resolved Rudder Model

The rudder is parameterised by the area scale s_A, the span scale s_H and the thickness ratio $\tau = t/c$ of a NACA four-digit symmetric section. The span, movable area, total area including the horn, geometric aspect ratio, mean chord and slipstream coverage become

$$\begin{aligned} H_R &= s_H H_{R0}, \qquad A_R = s_A A_{R0}, \qquad A_T = s_A A_{T0}, \\ \Lambda &= \frac{H_R^2}{A_T} = \Lambda_0 \frac{s_H^2}{s_A}, \qquad \bar{c} = \frac{A_T}{H_R}, \qquad \eta = \frac{D_P}{H_R}, \end{aligned} \tag{5}$$

With $\Lambda_0 = 1.827$ and $A_{T0} = H_{R0}^2/\Lambda_0$. Instead of replacing the measured lift-gradient coefficient $f_{\alpha 0} = 2.747$ by an empirical value, the model scales it relative to the baseline:

$$f_\alpha(\Lambda, \tau) = f_{\alpha 0} \, \frac{C_{N\infty}(\tau)}{C_{N\infty}(\tau_0)} \, \frac{\Lambda/(\Lambda+2.25)}{\Lambda_0/(\Lambda_0+2.25)} \tag{6}$$

Where the aspect-ratio factor is Fujii's law [12] and $C_{N\infty}(\tau)$ is the RANS-derived normal-force slope of NACA sections at infinite aspect ratio [10], interpolated over $\tau \in [0.12, 0.25]$; $\tau_0 = 0.18$ is the thickness of the original NACA 0018 rudder [10]. Equation (6) returns exactly the captive-test value for the original rudder, so the validated manoeuvring behaviour of the baseline is preserved, and it propagates changes of planform and section through physically grounded ratios.

The linear normal-force law is supplemented by a stall model. Behind a working propeller the stall angle of a rudder increases to roughly 30–$35°$ for aspect ratios around 2.6 [6, 10]; the KVLCC2 captive tests show a linear normal force up to $35°$ [13]. The baseline stall angle in the slipstream is therefore set to $\alpha_{s0} = 37°$ and shifted for a new design by the change of the two-dimensional section stall angle and by the lifting-line change of the induced angle at stall:

$$\begin{aligned} \alpha_s(\Lambda, \tau) &= \alpha_{s0} + [\alpha_s^{2D}(\tau, Re_R) - \alpha_s^{2D}(\tau_0, Re_R)] + \frac{C_{N,s0}}{\pi}\left(\frac{1}{\Lambda} - \frac{1}{\Lambda_0}\right) \\ C_{N,s0} &= f_{\alpha 0} \sin\alpha_{s0}. \end{aligned} \tag{7}$$

Beyond stall the normal-force coefficient decays smoothly towards a fraction $1 - k_p$ of its stall value,

$$C_N(\alpha) = \mathrm{sgn}(\alpha) \begin{cases} f_\alpha \sin|\alpha|, & |\alpha| \leq \alpha_s, \\ f_\alpha \sin\alpha_s \left[1 - k_p\left(1 - e^{-(|\alpha| - \alpha_s)/\Delta_p}\right)\right], & |\alpha| > \alpha_s, \end{cases} \tag{8}$$

With $k_p = 0.25$ and $\Delta_p = 4°$; Sect. 7.8 examines the sensitivity to α_{s0} and k_p.

Section data are obtained with NeuralFoil [15], a physics-informed neural surrogate of XFoil that covers Reynolds numbers up to 10^{10}. Transition is forced at 2% chord because the rudder operates in the turbulent propeller slipstream. Tables of the zero-lift drag $C_{D0}(\tau, Re)$, the two-dimensional stall angle and the minimum pressure coefficient $C_{p,\min}(\alpha, \tau)$ were generated for $\tau \in [0.10, 0.26]$ and $Re \in [2 \times 10^7, 1.2 \times 10^8]$. The profile drag of the rudder in the slipstream on a straight course is

$$R_{app} = 1/2\,\rho U_{R0}^2 A_T C_{D0}(\tau, Re_R), \qquad Re_R = U_{R0}\bar{c}/\nu \tag{9}$$

Where U_{R0} follows from Eq. (4); the baseline value is part of the measured resistance R'_0, so only the difference $\Delta R_{app} = R_{app} - R_{app,0}$ enters the surge balance of a modified design. The cavitation-inception angle $\alpha_{cav}(\tau)$ is the smallest angle of attack at which $-C_{p,\min}$ reaches the cavitation number $\sigma_R = (p_{atm} + \rho g h_c - p_v)/(1/2\,\rho U_s^2)$, evaluated at the depth h_c of the upper edge of the propeller race with the fully developed slipstream velocity U_s.

3.3. Environmental Loads

Wind loads follow Blendermann's parametric model [32] with the coefficients for a loaded tanker (lateral drag $C_{Dt} = 0.70$, longitudinal drag $0.90/0.55$ for head/stern wind, cross-force parameter $\delta_B = 0.40$) and an above-water geometry representative of a loaded VLCC ($A_F = 1400$ m², $A_L = 3500$ m², centroid 24 m aft of midship). Wind fluctuations are described by the NPD (Frøya) gust spectrum [33],

$$S_u(f) = \frac{320\,(U_{10}/10)^2}{\left(1+\tilde{f}^{\,n}\right)^{5/(3n)}}, \qquad \tilde{f} = 172 f (U_{10}/10)^{-3/4}, \quad n = 0.468 \tag{10}$$

And enter the loads through the derivative of the wind model with respect to the wind speed.

Waves are long-crested and described by the JONSWAP spectrum $S_\zeta(\omega)$ with $\gamma = 3.3$ [34]. The mean wave-drift loads follow the short-wave asymptotic waterline integral of Faltinsen et al. [21] with the finite-draft correction of Fujii and Takahashi [22]:

$$\bar{F}_i = \rho g\, I_i(\chi) \int_0^\infty S_\zeta(\omega)\, \alpha_T(\omega)\, d\omega, \qquad I_i(\chi) = \int_{\Gamma_\chi} \left(\hat{\mathbf{k}} \cdot \hat{\mathbf{n}}\right)^2 (-n_i)\, dl \tag{11}$$

$$\alpha_T = \frac{\pi^2 I_1^2(1.5kd)}{\pi^2 I_1^2(1.5kd) + K_1^2(1.5kd)} \tag{12}$$

Where $i \in \{x, y, \psi\}$ with $n_\psi = x n_y - y n_x$, Γ_χ is the part of the waterline illuminated by waves from the relative direction χ ($\chi = 0$ head seas), $\hat{\mathbf{k}}$ is the propagation direction and I_1, K_1 are modified Bessel functions. The design waterline is represented parametrically (parallel middle body $0.5L_{pp}$, elliptic bow $0.2L_{pp}$, power-law stern $0.3L_{pp}$; waterplane coefficient 0.871). Because the transfer function of Eq. (11) factorises into a geometric and a spectral part, Newman's approximation [17] gives the slowly varying drift loads in closed form,

$$S_{F_i F_j}(\mu) = I_i I_j\, \Phi(\mu), \qquad \Phi(\mu) = 8 \int_0^\infty S_\zeta(\omega) S_\zeta(\omega+\mu) [1/2\,\rho g\, \alpha_T(\omega + \mu/2)]^2 d\omega \tag{13}$$

First-order sway force and yaw moment are computed with a long-wave slender-body strip approximation in which each section experiences the Froude–Krylov and added-mass loads $[\rho A_s(x) + a_{22}(x)]\, a_y(x,t)$, with a_y the transverse fluid acceleration at mid-draft, the sectional area $A_s(x)$ derived from the waterline and $a_{22}(x)$ distributed so that it integrates to the measured sway added mass m_y. These loads act at the encounter frequency $\omega_e = \omega + kU\cos\chi$.

3.4. Operational Climate From Buoy Records

The operational climate is derived from NOAA National Data Buoy Center records [18]. The joint distribution of significant wave height and energy period comes from 194,475 records of station 46050 (Stonewall Bank, off Oregon); peak periods follow from $T_e/T_p = 0.903$ for the JONSWAP shape. The marginal distribution of the wind speed comes from 28,468 ten-minute records of offshore station 46002 (2016), and the dependence between wind speed and wave height is estimated from

2,164 concurrent records of station 46097 (Spearman $\rho_S = 0.083$), which reflects the swell-dominated wave climate of the North-East Pacific. A Gaussian copula with parameter $2\sin(\pi\rho_S/6)$ couples the marginals, and 6×10^4 samples are reduced by k-means clustering to eight sea states (Table 3). Wind and waves are taken as co-directional and the relative heading is uniformly distributed; with port–starboard symmetry seven headings $\chi \in \{0°, 30°, \ldots, 180°\}$ represent the circle. The service climate thus comprises $C = 56$ weighted conditions.

3.5. Heading Autopilot And Steering Gear

The course-keeping controller is a PID autopilot acting on low-pass-filtered heading and yaw-rate measurements:

$$\delta_c = -K_p \left[\psi_f + T_d r_f + \frac{1}{T_I}\int_0^t \psi_f \, dt'\right], \qquad \dot{\psi}_f = \frac{\psi - \psi_f}{T_f}, \quad \dot{r}_f = \frac{r - r_f}{T_f} \tag{14}$$

Where ψ is the heading deviation from the set course, K_p the gain, T_d the derivative time, T_f the filter time constant and $T_I = 400$ s the integral time. The filter prevents the autopilot from reacting to wave-frequency yaw motions. The steering gear is modelled as $\dot{\delta} = \mathrm{sat}_{\dot{\delta}_{\max}}[(\delta_c - \delta)/T_E]$ with $T_E = 2.5$ s. K_p, T_d and T_f are design variables.

4. Spectral evaluation of rudder-related power and course-keeping precision

4.1. Steady sailing equilibrium

For each design and condition c (sea state, wind and relative heading), the mean state is obtained from the MMG equations with all accelerations set to zero and the mean environmental loads of Sect. 3.3. Two modes are used. In the speed-maintaining mode, which represents service operation, $u = U_0$ is prescribed and (v_0, δ_0, n_0) solve $X = Y = N = 0$. In the fixed-revolution mode, which represents adverse-weather operation at the installed service revolutions, n is prescribed and (u_0, v_0, δ_0) are solved for. Newton's method with a finite-difference Jacobian and step limitation converges in a few iterations; conditions without an equilibrium are flagged infeasible.

4.2. Closed-loop linear model

Linearising the sway and yaw equations of Eq. (1) about the equilibrium at constant forward speed, and appending heading kinematics, steering gear, measurement filters and integrator, gives the closed-loop model

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{b}_g u_g + \mathbf{B}_d \mathbf{d}_{LF} + \mathbf{B}_d \mathbf{d}_{WF}, \qquad \mathbf{x} = \left[v, r, \psi, \delta, \psi_f, r_f, z\right]^{\mathrm{T}} \tag{15}$$

Where $z = \int \psi_f \, dt$, u_g is the gust velocity, $\mathbf{d}_{LF} = [Y, N]^{\mathrm{T}}_{sv}$ the slowly varying drift loads and $\mathbf{d}_{WF} = [Y, N]^{\mathrm{T}}_{1}$ the first-order wave loads. The upper rows of $\mathbf{A}$ contain $\mathbf{M}^{-1}$ times the derivatives of the total sway force and yaw moment with respect to (v, r, ψ, δ), including the Coriolis terms of Eq. (1) and the dependence of the wind and drift loads on the heading; $\mathbf{M}$ is the sway–yaw mass matrix. All derivatives are central finite differences of the full nonlinear force model, so the asymmetric flow-straightening and wake coefficients and the stall model are linearised consistently about the actual operating point. Designs for which $\mathbf{A}$ has an eigenvalue with non-negative real part in any condition are infeasible.

4.3. Frequency-domain Second Moments

With the eigendecomposition $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$, the frequency response $\mathbf{G}(i\omega) = \mathbf{V}(i\omega\mathbf{I} - \mathbf{\Lambda})^{-1}\mathbf{V}^{-1}$ is evaluated for all frequencies at the cost of one eigendecomposition per case. The stationary covariance of the state is the sum of three independent contributions,

$$\begin{aligned}\boldsymbol{\Sigma} = \quad & \sum_k \mathrm{Re}\left[\mathbf{G}\mathbf{b}_g\mathbf{b}_g^{\mathrm{H}}\mathbf{G}^{\mathrm{H}}\right]_{\omega_k} S_u(\omega_k)\Delta\omega_k \\ & + \sum_k \mathrm{Re}\left[\mathbf{G}\mathbf{B}_d\mathbf{I}\mathbf{I}^{\mathrm{T}}\mathbf{B}_d^{\mathrm{T}}\mathbf{G}^{\mathrm{H}}\right]_{\mu_k} \Phi(\mu_k)\Delta\mu_k \\ & + \sum_k \mathrm{Re}\left[\mathbf{G}\mathbf{B}_d\mathbf{h}_k\mathbf{h}_k^{\mathrm{H}}\mathbf{B}_d^{\mathrm{T}}\mathbf{G}^{\mathrm{H}}\right]_{\omega_{e,k}} S_\zeta(\omega_k)\Delta\omega_k,\end{aligned} \tag{16}$$

Where $\mathbf{I} = \left[I_y, I_\psi\right]^{\mathrm{T}}$ is the geometric drift vector of Eq. (13) and $\mathbf{h}_k$ the complex first-order load vector per unit wave amplitude. Ninety logarithmically spaced frequencies between 5×10^{-4} and 0.4 rad/s resolve the low-frequency part and 64 wave components between 0.2 and 1.8 rad/s the wave-frequency part. The standard deviations σ_ψ and σ_δ and the rate demand $\sigma_{\dot\delta} = \sqrt{\mathbf{c}^{\mathrm{T}}\boldsymbol{\Sigma}\mathbf{c}}$, with $\mathbf{c}$ the row of $\mathbf{A}$ that defines $\dot\delta$, follow directly.

4.4. Mean Steering Resistance From The Surge Balance

At constant speed the propeller must deliver the effective thrust that balances the surge force

$$F_x(v, r, \delta) = X_H + X_R + \left(m + m_y\right)vr + x_G m r^2 \tag{17}$$

Which collects hull, rudder and inertial terms of Eq. (1); environmental surge loads are independent of the rudder system and are excluded. Expanding F_x to second order about the equilibrium $\mathbf{s}_0 = (v_0, 0, \delta_0)$ and taking expectations over the zero-mean fluctuations $\tilde{\mathbf{s}} = \left(\tilde v, \tilde r, \tilde\delta\right)$ yields

$$\mathbb{E}[-F_x] = -F_x(\mathbf{s}_0) - 1/2\ \mathrm{tr}(\mathbf{H}_x\boldsymbol{\Sigma}_s) + \mathcal{O}(\parallel \tilde{\mathbf{s}} \parallel^3) \tag{18}$$

With $\mathbf{H}_x$ the Hessian of F_x with respect to $\mathbf{s}$ (finite differences) and $\boldsymbol{\Sigma}_s$ the corresponding block of Eq. (16). The steering resistance in condition c therefore splits into a weather-helm part $\Delta R_{st,c} = -[F_x(\mathbf{s}_0) - F_x(\mathbf{0})]$ and a fluctuation part $\Delta R_{fl,c} = -1/2\,\mathrm{tr}(\mathbf{H}_x\boldsymbol{\Sigma}_s)$. Equation (18) generalises the quadratic steering criteria of Koyama [8] and Reid [9]: the weights of the yaw, sway and rudder variances are not prescribed but derived from the hydrodynamic model of the specific design and operating point, and cross-covariances such as $\mathbb{E}\left[\tilde v\tilde\delta\right]$, which capture whether the rudder moves with or against the local cross-flow, are retained.

4.5. Design metrics

The expected rudder-related effective power over the climate is

$$f_E = U_0\left[R_{app} + \sum_{c=1}^{C} w_c\left(\Delta R_{st,c} + \Delta R_{fl,c}\right)\right] \tag{19}$$

And the course-keeping precision is the climate-averaged root-mean-square heading deviation $f_\psi = \left(\sum_c w_c\,\sigma_{\psi,c}^2\right)^{1/2}$. Three further metrics enter the constraints. The adverse-weather rudder demand $\delta_{adv} = \max_c\left(\left|\delta_{0,c}\right| + 2\sigma_{\delta,c}\right)$ is evaluated in fixed-revolution mode for the adverse conditions of MEPC.1/Circ.850/Rev.3 for ships longer than 250 m ($h_s = 6.0$ m, $T_p \in \{7,9,11,13,15\}$ s, $V_w = 22.6$ m/s) with wind and waves from seven directions between head and following seas [2]. The steering-gear rate demand is $2\max_c\sigma_{\dot\delta,c}$ over the service conditions. The service rudder-angle quantile δ_{95} is the 95% quantile of $|\delta|$ under the Gaussian mixture $\sum_c w_c\,\mathcal{N}\left(\delta_{0,c}, \sigma_{\delta,c}^2\right)$ and is compared with the cavitation-inception angle.

5. Multi-objective co-design problem

5.1. Decision Variables, Objectives And Constraints

The co-design vector $\mathbf{x} = (s_A, s_H, \tau, K_p, T_d, T_f)$ combines the rudder planform and section with the autopilot parameters (Table 1). The span may exceed the original value by at most 5%, the margin assumed to be available between rudder sole and baseline, and the area range covers rudders from half to 1.4 times the original size. The problem reads

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{f}(\mathbf{x}) = \left[f_E(\mathbf{x}),\ U_{IMO}(\mathbf{x}),\ f_\psi(\mathbf{x}) \right] \\ \text{s.t.} \quad & g_1 = U_{IMO} - 1 \le 0, \qquad g_2 = \delta_{adv} - 25^\circ \le 0, \qquad g_3 = 2\max_c \sigma_{\dot{\delta},c} - \dot{\delta}_{\max} \le 0 \\ & g_4 = \delta_{95} - \alpha_{cav} \le 0, \qquad g_5\colon \text{ equilibria exist and } \max \mathrm{Re}\, \lambda(\mathbf{A}_c) < 0 \ \ \forall c. \end{aligned} \tag{20}$$

The second objective is the utilisation of the IMO manoeuvring criteria [1],

$$U_{IMO} = \max_{\pm} \max \left\{ \frac{A_D}{4.5L}, \frac{T_D}{5L}, \frac{d_{10}}{2.5L}, \frac{\theta_1^{10/10}}{20^\circ}, \frac{\theta_2^{10/10}}{40^\circ}, \frac{\theta_1^{20/20}}{25^\circ} \right\} \tag{21}$$

Where A_D and T_D are the advance and tactical diameter of the 35° turning circle, d_{10} the track length travelled until the heading has changed by 10° with 10° rudder (initial turning ability), and θ the overshoot angles of the zig-zag manoeuvres; the overshoot limits correspond to $L/V = 40.1$ s. The maximum is taken over port and starboard manoeuvres, so six simulations are performed per design. $U_{IMO} < 1$ indicates compliance and $1 - U_{IMO}$ is the safety margin of the most critical criterion. The rudder-demand limit g_2 keeps a reserve of 10° to the hard-over angle for manoeuvring while the ship keeps course in the adverse conditions; g_3 ensures that the steering gear can follow the autopilot without frequent rate saturation; g_4 keeps the rudder free of sheet cavitation for 95% of the service time.

Table 1. Decision variables and bounds. The reference design B0 combines the original KVLCC2 rudder with a conventionally tuned autopilot (Sect. 5.4).

Variable	Meaning	Lower	Upper	B0
s_A	rudder area scale (A_R/A_{R0})	0.50	1.40	1.00
s_H	rudder span scale (H_R/H_{R0})	0.80	1.05	1.00
τ	section thickness ratio t/c (NACA 00xx)	0.12	0.25	0.18
K_p	autopilot gain (–)	0.5	5.0	2.76
T_d	derivative time (s)	10	150	78.5
T_f	measurement-filter time constant (s)	1	40	40.0

5.2. Vectorised Evaluation

A design evaluation consists of the six IMO manoeuvres in the time domain (fourth-order Runge–Kutta, $\Delta t = 1$ s at full scale) and $C + 35 = 91$ spectral cases (56 service and 35 adverse conditions). All parts are vectorised over candidates: the manoeuvres of an entire population are integrated simultaneously, and equilibria, linearisations and frequency responses are computed for all design–condition pairs in batched linear algebra. One complete evaluation costs about 38 ms on a single CPU core, which makes populations of 64 designs over 80 generations feasible in about four minutes.

5.3. Robust Formulation

The captive-model coefficients carry measurement and analysis uncertainty, and benchmark comparisons between facilities show appreciable scatter [14]. The robust variant treats the

hydrodynamic coefficients as independent Gaussian random multipliers with the coefficients of variation listed in Table 2, and replaces the second objective and the adverse-weather constraint by their 90% quantiles over $S = 12$ common random samples:

$$U_{IMO}^{0.9}(\mathbf{x}) = Q_{0.9}\{U_{IMO}(\mathbf{x}, \boldsymbol{\xi}_s)\}_{s=1}^{S}, \qquad \delta_{adv}^{0.9}(\mathbf{x}) = Q_{0.9}\{\delta_{adv}(\mathbf{x}, \boldsymbol{\xi}_s)\}_{s=1}^{S} \le 25^\circ \tag{22}$$

Common random numbers make the comparison between candidates consistent; the robustness of the resulting designs is verified afterwards with an independent sample of 200 realisations.

Table 2. Coefficients of variation of the hydrodynamic parameters used in the robust formulation and in the Monte Carlo verification

Parameter group	CoV
Linear hull derivatives Y'_v, Y'_r, N'_v, N'_r	0.05
Nonlinear hull derivatives ($Y'_{vvv}, \dots, N'_{rrr}$) and surge derivatives	0.15
Hull–rudder interaction t_R, a_H, x'_H; flow straightening $\gamma_R^{\pm}$, l'_R; ε, κ	0.10
Wake and thrust deduction w_{P0}, t_P; lift gradient $f_{\alpha 0}$	0.05

5.4. Reference Designs And Competing Strategies

The following strategies are compared at an equal budget of 5120 evaluations:

B0 (reference). The original KVLCC2 rudder with an autopilot tuned by minimising the Koyama-type criterion $\mathbb{E}[\psi^2] + \lambda\mathbb{E}[\delta^2]$, $\lambda = 0.1$, over the climate, subject to the same constraints g_2–g_5 (differential evolution, 1664 evaluations). This represents the conventional practice of tuning the controller for an installed rudder.

Fixed-AP. Geometry-only optimisation (s_A, s_H, τ) with the autopilot of B0.

Sequential. Geometry optimisation with the B0 autopilot using half of the budget, followed by autopilot optimisation (K_p, T_d, T_f) for eight geometries spread along the stage-one front using the other half.

Co-design (proposed). Simultaneous optimisation of all six variables.

Co-design with original section. As above with the thickness ratio fixed at $\tau = 0.18$, which isolates the contribution of planform and autopilot.

Robust co-design. Co-design of all six variables with the quantile formulation of Eq. (22).

The co-design problem is solved with NSGA-II [23]; NSGA-III [24], SMS-EMOA [25], AGE-MOEA [26] and Latin-hypercube random search with the same budget serve as algorithmic baselines. All algorithms use the pymoo implementations [27] with population size 64 (66 reference directions for NSGA-III), Latin-hypercube initialisation, simulated binary crossover and polynomial mutation with default settings, and constraint-domination for feasibility handling.

6. Experimental Setup

6.1. Data

The study uses three groups of real data: (i) the KVLCC2 captive-model coefficients and propeller open-water data of the MMG standard method [13, 20]; (ii) the MARIN free-running model-test results for the $\pm 35^\circ$ turning circles and the $\pm 10^\circ/10^\circ$ and $\pm 20^\circ/20^\circ$ zig-zag manoeuvres [19, 20], used for validation; and (iii) NOAA NDBC buoy records of stations 46050, 46002 and 46097 [18], used for the climate. Section data come from NeuralFoil [15] and the wind-load coefficients from Blendermann [32].

6.2. Metrics

Front quality is measured with the hypervolume (HV) [28] of the feasible non-dominated set among all designs evaluated by a run. Objectives are normalised with the ideal and nadir points of the union of all feasible fronts of all runs, and the reference point is 1.1 in each normalised objective. Convergence is reported as HV versus the number of evaluations. The knee design of a front is the feasible non-dominated design closest to the ideal point in normalised objective space. The co-design, the sequential strategy and the algorithmic baselines are run with three random seeds; the fixed-autopilot, fixed-section, ablation and robust variants with one seed each.

6.3. Verification and Sensitivity Protocol

The spectral evaluator is verified against nonlinear time-domain simulation (Sect. 7.3). Global sensitivity is quantified with Sobol indices [30] estimated by Saltelli sampling [31] with SALib [35] ($N = 256$, 3584 evaluations). Modelling assumptions are varied one at a time for selected designs, and the robustness of nominal and robust Pareto designs is verified by Monte Carlo simulation with 200 independent realisations of the hydrodynamic coefficients.

6.4. Implementation

The framework is implemented in Python with NumPy/SciPy and pymoo; all computations ran on a single CPU core. The code, data-processing scripts and results are provided with the paper.

7. Results

7.1. Operational climate

Figure 2 shows the buoy data and the eight sea states, and Table 3 lists them. The climate is dominated by moderate seas: 66.6% of the time falls into the four clusters with $H_s < 2$ m, and the most severe cluster ($H_s = 5.65$ m, $T_p = 12.8$ s) occurs 5.4% of the time. Because the wind–wave dependence at the coast of the North-East Pacific is weak, wind speeds of 3–12 m/s occur across all wave clusters.

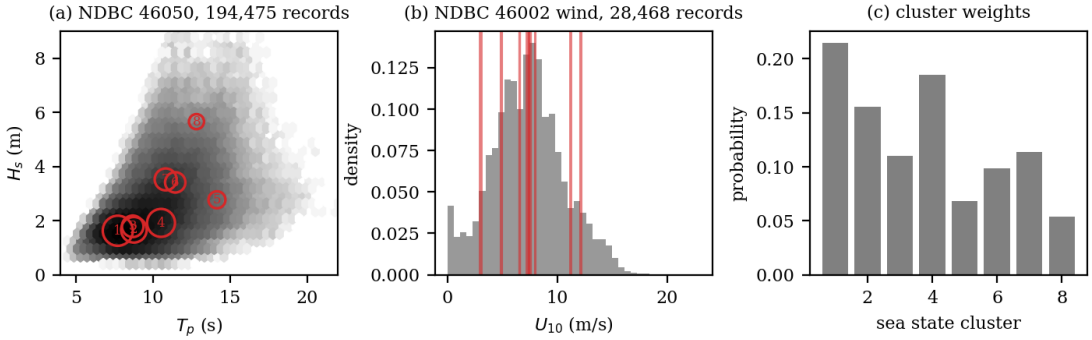


Figure 2. Operational climate from NOAA NDBC records: (a) joint distribution of peak period and significant wave height at station 46050 with the eight cluster centroids (circle area proportional to probability); (b) wind-speed distribution at station 46002 with the cluster wind speeds; (c) cluster probabilities.

Table 3. Representative sea states obtained by k-means clustering of 6×10^4 copula samples

Cluster	1	2	3	4	5	6	7	8
H_s (m)	1.63	1.66	1.79	1.92	2.77	3.41	3.52	5.65
T_p (s)	7.72	8.76	8.69	10.52	14.15	11.45	10.82	12.82
U_{10} (m/s)	7.34	2.99	12.15	7.46	6.51	11.20	4.92	7.92
Probability	0.215	0.156	0.110	0.185	0.068	0.098	0.114	0.054

7.2. Validation of the Manoeuvring Model

Table 4 compares the standard manoeuvres of the original KVLCC2 predicted by the present implementation with the MARIN free-running model tests and with the predictions of Yasukawa and Yoshimura [13] that use the same coefficients; Fig. 3 shows the trajectories and time histories. The present implementation reproduces the reference MMG predictions closely: the mean absolute difference is $0.21L$ for the turning indices and $0.76°$ for the overshoot angles. With respect to the free-running tests, the mean absolute error of the turning indices is $0.11L$ (reference MMG prediction: $0.10L$), and the overshoot angles show the same tendency to underestimation as the reference MMG prediction (mean absolute error $3.96°$ versus $3.20°$), which is a known property of this coefficient set [10, 20]. The ranking of manoeuvres and the port–starboard asymmetry are reproduced, which is what matters for the relative comparison of rudder designs. Halving the time step from 1 s to 0.5 s changes U_{IMO} by 0.3%.

Table 4. Standard manoeuvres of the original KVLCC2: MARIN free-running model tests, MMG prediction of Yasukawa and Yoshimura (Y&Y) [13] as tabulated in [20], and present implementation

Index	MARIN	Y&Y	Present	Present vs MARIN
Advance, $+35°$ (L)	3.25	3.31	3.07	-5.5%
Tactical diameter, $+35°$ (L)	3.34	3.36	3.11	-6.9%
Advance, $-35°$ (L)	3.11	3.26	3.08	-0.8%
Tactical diameter, $-35°$ (L)	3.08	3.26	3.10	$+0.5\%$
1st overshoot, $+10°/10°$ (deg)	8.2	5.2	5.0	-39.6%
2nd overshoot, $+10°/10°$ (deg)	21.9	15.8	13.9	-36.5%
1st overshoot, $-10°/10°$ (deg)	9.5	7.6	7.3	-23.7%
2nd overshoot, $-10°/10°$ (deg)	15.0	10.2	9.0	-40.1%
1st overshoot, $+20°/20°$ (deg)	13.7	10.9	10.6	-22.6%
1st overshoot, $-20°/20°$ (deg)	15.1	14.5	14.0	-7.6%

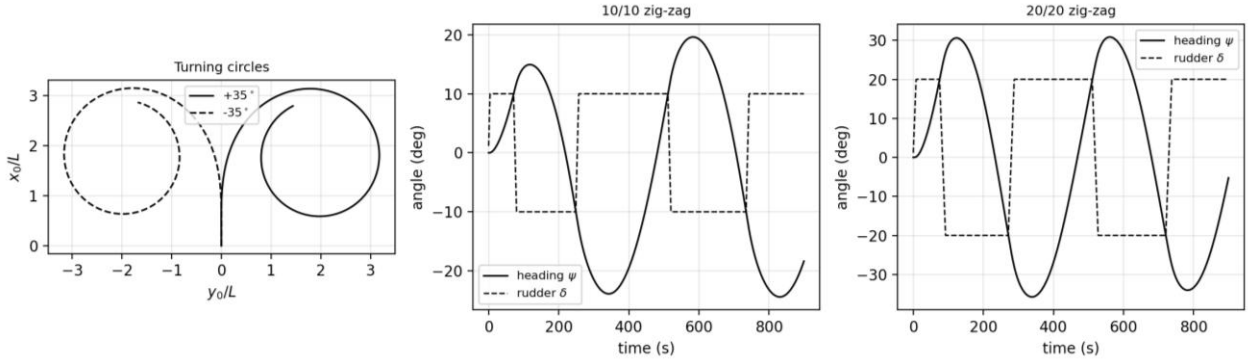


Figure 3. Simulated standard manoeuvres of the original KVLCC2 at 15.5 kn: $\pm 35°$ turning circles and $+10°/10°$ and $+20°/20°$ zig-zag manoeuvres

7.3. Verification of the Spectral Evaluator

The spectral evaluator was verified against nonlinear time-domain simulations that use the full MMG sway and yaw forces at constant speed, the rate-limited steering gear, the nonlinear Blendermann wind loads with synthesised NPD gusts, heading-dependent slowly varying drift loads generated with Newman's approximation in the time domain and first-order wave loads synthesised from the same JONSWAP spectra. Six designs, including B0 and designs with area scales between 0.6 and 1.25 and a wide range of autopilot settings, were simulated in the eight sea states at relative headings of $45°$ and $135°$, i.e. 96 cases. Each case was simulated for 16,384 s (4.6 h, the first 1500 s discarded) with two independent random seeds.

Figure 4 and Table 5 summarise the agreement. The heading and rudder standard deviations agree with $R^2 = 0.996$ and median relative errors of 3.0% and 4.8%. The steering resistance agrees with $R^2 = 0.976$ and a median relative error of 10.9%, which equals the statistical uncertainty of the time-domain reference itself (seed-to-seed scatter 11.3%). The mean rudder angle carries a small systematic offset ($R^2 = 0.948$) because nonlinear sway–yaw coupling shifts the mean state when fluctuations are large. Most importantly for optimisation, the ranking of the six designs by mean rudder-related power is identical in both methods (Spearman correlation 1.0). Rate saturation occurred in 3 of the 96 cases, all in the most severe sea state, and the agreement is unchanged when these cases are excluded. The spectral evaluator needs 0.75 ms per case against 0.84 s for a time-domain record of 4.6 h, a speed-up of about 1100; a complete design evaluation over the 56-condition climate would otherwise require about 47 s of simulation per design and per random seed, which would make population-based optimisation impractical.

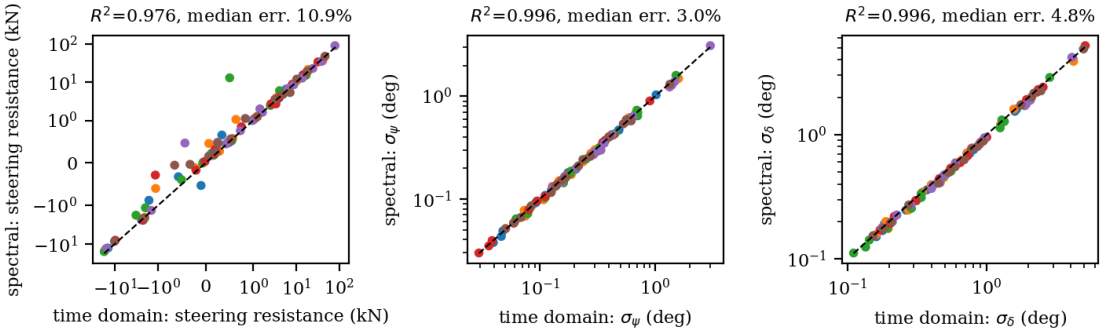


Figure 4. Spectral evaluator versus nonlinear time-domain simulation for 96 cases (6 designs, colour-coded; 8 sea states; 2 headings): steering resistance, heading standard deviation and rudder-angle standard deviation; time-domain values are averages over two seeds

Table 5. Agreement between the spectral evaluator and nonlinear time-domain simulation (96 cases)

Quantity	R^2	Pearson r	Median relative error
Steering resistance $\Delta R_{st} + \Delta R_{fl}$	0.976	0.989	10.9%
Heading standard deviation σ_ψ	0.996	0.998	3.0%
Rudder standard deviation σ_δ	0.996	0.999	4.8%
Mean rudder angle δ_0	0.948	0.998	12.5%
Time-domain seed-to-seed scatter: 11.3% (resistance), 3.6% (σ_ψ)			
Computing time per case: 0.75 ms (spectral) vs 0.84 s (time domain), speed-up ≈ 1100			

7.4. Pareto-optimal Rudder–Autopilot Co-Designs

The three NSGA-II runs of the co-design problem were merged into a front of 1244 feasible non-dominated designs (Fig. 5). The front spans rudder-related powers from 216.9 to 336.9 kW, IMO utilisations from 0.648 to 0.770 and heading deviations from 0.203° to 0.543°. Its structure is informative. In the f_E–U_{IMO} plane the designs collapse onto a narrow curve: energy and calm-water manoeuvring margin are traded through the rudder area, whose correlation with f_E along the front is 0.999 (Fig. 6). The heading deviation adds an almost independent dimension that is controlled by the autopilot: for a given area, the gain and derivative time move the design between precise and economical course keeping. The reference design B0 lies clearly inside the attainable region; 187 front designs dominate it in all three objectives.

Table 6 lists B0 and five representative co-designs: the knee of the front, the most energy-efficient design that still dominates B0 (B0-dom.), and the three extreme designs. The knee design reduces the rudder-related power from 315.2 to 267.8 kW (−15.0%) while improving the IMO utilisation from

0.705 to 0.691 and the heading deviation from $0.255°$ to $0.240°$; its adverse-weather rudder demand (20.0°) and steering-gear rate demand (2.15°/s) are also lower than those of B0. The B0-dominating design with the lowest power saves 18.9% with 10% less rudder area. At the low-energy end, the minimum-power design uses 61% of the original area and an aspect ratio of 3.28; it saves 31.2%, but its heading deviation nearly doubles, its adverse-weather demand reaches the 25° limit, and with a stall angle of 28.9° it is the only selected design that stalls during the $35°$ turning manoeuvre. At the safe end, rudders with 1.4 times the original area reach $U_{IMO} = 0.648$ at a power penalty of 5–7%.

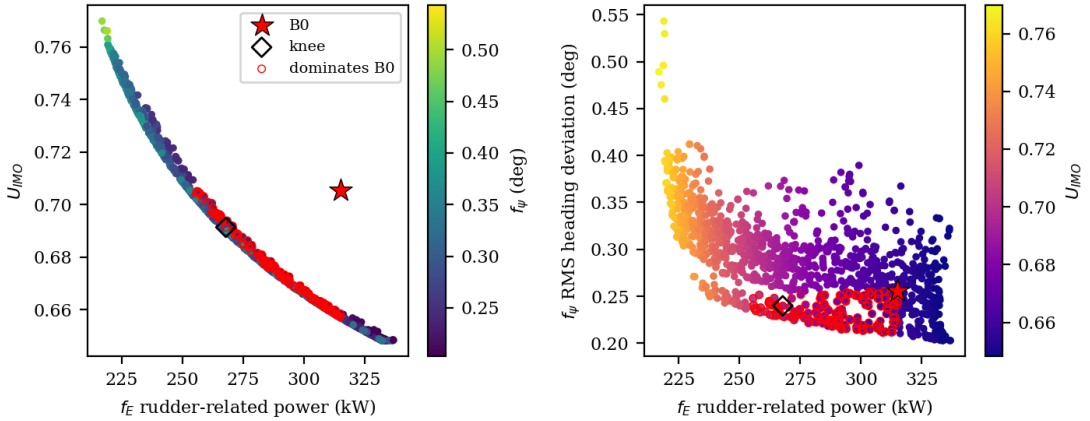


Figure 5. Merged Pareto front of the co-design problem (three NSGA-II runs, 1244 designs): rudder-related power versus IMO utilisation (colour: heading deviation) and versus heading deviation (colour: IMO utilisation). The star marks B0, the diamond the knee design and red circles the designs that dominate B0

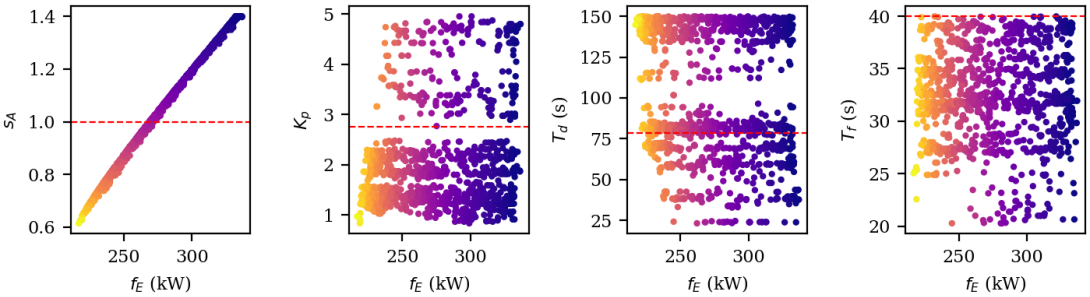


Figure 6. Rudder area scale and autopilot parameters along the Pareto front (colour: U_{IMO}, scale as in Fig. 5b); dashed lines mark B0. The span scale (1.04–1.05) and thickness ratio (0.120–0.125) are nearly constant along the front and are not shown

Table 6. Reference design B0 and selected Pareto-optimal co-designs. Manoeuvring indices are the worse of port and starboard

	B0	Knee	B0-dom.	Min-f_E	Min-U_{IMO}	Min-f_ψ
s_A	1.000	0.985	0.901	0.614	1.400	1.399
s_H	1.000	1.050	1.050	1.050	1.050	1.050
τ	0.180	0.120	0.120	0.120	0.120	0.124
K_p	2.76	3.35	3.39	0.97	0.91	4.80
T_d (s)	78.5	61.1	58.3	144.8	144.8	37.7
T_f (s)	40.0	39.0	37.4	25.0	30.1	31.9
Λ	1.83	2.05	2.24	3.28	1.44	1.44
f_α	2.75	3.03	3.17	3.77	2.48	2.48
α_s (deg)	37.0	34.5	33.2	28.9	41.0	41.0
$C_{D0} \times 10^3$	6.64	5.82	5.90	6.26	5.52	5.57
Profile drag power (kW)	265.0	221.8	205.7	148.7	299.1	301.6
Steering-induced power (kW)	50.2	46.1	49.9	68.2	33.2	35.3
f_E (kW)	315.2	267.8	255.6	216.9	332.3	336.9
Δf_E vs B0	+0.0%	-15.0%	-18.9%	-31.2%	+5.4%	+6.9%
U_{IMO}	0.705	0.691	0.704	0.770	0.648	0.648
f_ψ (deg)	0.255	0.240	0.248	0.489	0.333	0.203
δ_{adv} (deg)	21.1	20.0	21.1	25.0	15.2	17.8
Rate demand (deg/s)	2.22	2.15	2.17	2.29	1.80	2.34
δ_{95} / α_{cav} (deg)	3.9 / 11.0	3.7 / 9.5	3.9 / 9.5	4.7 / 9.5	3.0 / 9.5	3.2 / 9.6
A_D/L	3.08	3.02	3.07	3.42	2.86	2.86
T_D/L	3.11	3.06	3.12	3.38	2.89	2.89
d_{10}/L	1.76	1.73	1.76	1.92	1.62	1.62
$\theta_1^{10/10}$ (deg)	7.2	6.7	7.2	8.2	6.2	6.2
$\theta_2^{10/10}$ (deg)	13.8	12.8	13.5	18.1	10.9	10.9
$\theta_1^{20/20}$ (deg)	13.9	13.2	13.9	15.7	12.7	12.7

Figure 7 decomposes the rudder-related power of the selected designs, and Table 7 traces the path from B0 to the knee design one design change at a time. Profile drag dominates the budget for every design. Replacing the NACA 0018 by a NACA 0012 section lowers it by 35.8 kW and, through the higher normal-force slope, also reduces the weather helm. Extending the span by 5% saves another 6.9 kW and the small area change 2.3 kW; the co-designed autopilot then converts a marginal 0.7 kW into a lower heading deviation and a lower steering-gear rate demand. All intermediate designs remain better than B0 in both safety objectives, so the knee design is reached without any compromise in manoeuvring safety.

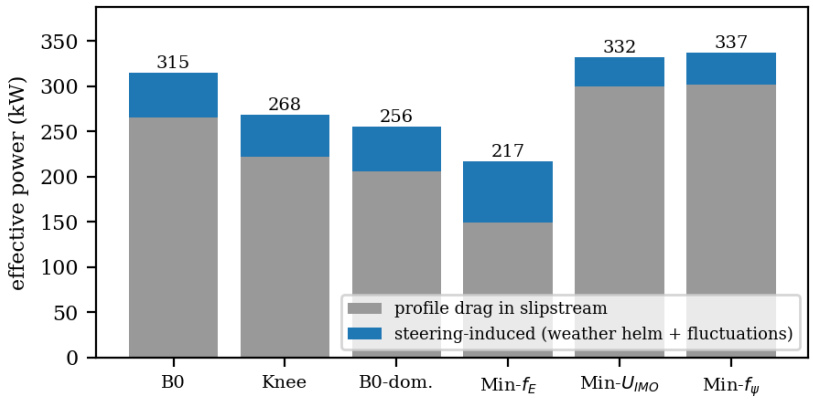


Figure 7. Rudder-related effective power of B0 and the selected co-designs, split into profile drag in the slipstream and steering-induced power (weather helm plus fluctuations)

Table 7. Step-by-step decomposition of the change from B0 to the knee design

Step	Profile drag	Steering	f_E	Δf_E	U_{IMO}	f_ψ	Rate demand
	(kW)	(kW)	(kW)	(kW)		(deg)	(deg/s)
B0	265.0	50.2	315.2	–	0.705	0.255	2.22
+ section τ: 0.18→0.12	229.2	47.1	276.3	-38.8	0.695	0.248	2.22
+ span s_H: 1→1.05	224.7	44.7	269.4	-6.9	0.689	0.243	2.22
+ area s_A: 1→0.985	221.8	45.3	267.1	-2.3	0.691	0.244	2.22
+ autopilot (= knee)	221.8	46.1	267.8	+0.7	0.691	0.240	2.15

The savings were finally checked with the nonlinear time-domain simulator over the complete service climate, i.e. all 56 conditions with 16,384 s each (Table 8). The time-domain steering-induced power is about 10 kW lower than the spectral estimate for every design, so the spectral evaluator is slightly conservative in absolute terms and the relative savings are preserved: 15.4% for the knee design (spectral 15.0%), 19.5% for the most efficient B0-dominating design (18.9%) and 33.0% for the minimum-power design (31.2%). The climate-averaged heading deviations agree within 1–6%.

Table 8. Climate-averaged rudder-related power and heading deviation of B0 and selected co-designs: spectral evaluator versus nonlinear time-domain simulation of all 56 service conditions (16,384 s per condition)

	f_E (kW)		Saving vs B0		f_ψ (deg)	
Design	spectral	time domain	spectral	time domain	spectral	time domain
B0	315.2	305.4	–	–	0.255	0.251
Knee	267.8	258.4	15.0%	15.4%	0.240	0.238
B0-dom.	255.6	245.8	18.9%	19.5%	0.248	0.245
Min-f_E	216.9	204.6	31.2%	33.0%	0.489	0.460

7.5. Co-design Versus Conventional Design Strategies

Table 9 and Fig. 8b compare the design strategies at the same budget of 5120 evaluations. Simultaneous co-design attains the largest hypervolume (1.0862 ± 0.0016), followed by the sequential strategy (1.0647 ± 0.0033) and the geometry-only optimisation with the autopilot of B0 (1.0565). The ranking is the same, and the differences are larger, for the B0-referenced hypervolume, which measures the volume of designs that improve on B0 in all three objectives: co-design reaches 3.01×10^{-3}, which is 1.29 times the value of the sequential strategy and 1.37 times that of the fixed autopilot. Every co-design run exceeds every sequential run, and every sequential run exceeds the fixed-autopilot run (Table 9), which is the most significant outcome attainable with three runs per

strategy (one-sided Mann–Whitney test, $p = 0.05$). Because they search only three variables, the geometry-only strategies converge faster at first, but co-design overtakes both after 1600 evaluations and stays ahead. The convergence curves (Fig. 8b) also show that the advantage is not a matter of budget: the second, autopilot-only stage of the sequential strategy adds only 0.8% to the hypervolume reached after its geometry stage, because geometries that are efficient only together with a different autopilot have already been discarded. The B0 autopilot also prevents the geometry-only optimisation from exploiting small rudders and precise course keeping: the lowest power it reaches is 227.2 kW against 216.9 kW for co-design, and its best heading deviation is $0.218°$ against $0.203°$. When the original NACA 0018 section is retained, co-design finds 12 designs that dominate B0, and the most efficient of them saves 3.6%, against 18.9% when the section is free. The thinner section therefore provides most of the attainable energy saving, consistent with the decomposition of Table 7, while span, area and autopilot add the remainder and secure the simultaneous gains in manoeuvring margin and heading precision.

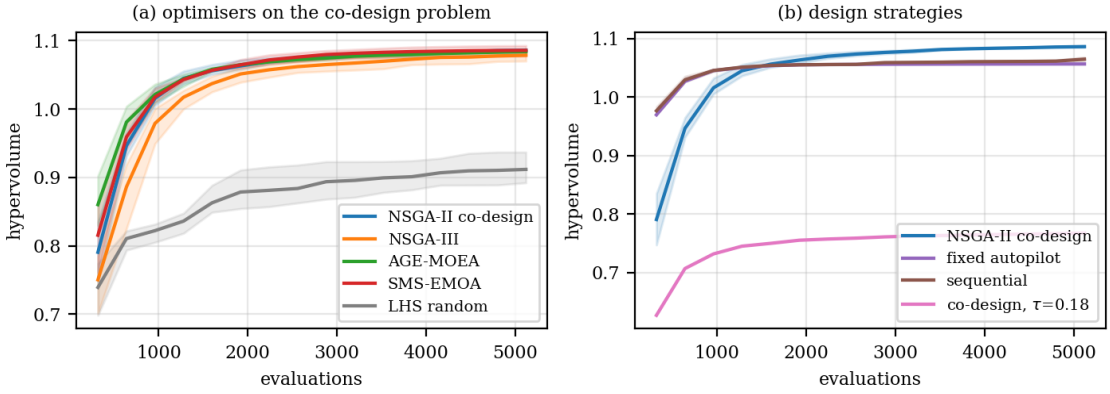


Figure 8. Hypervolume of the feasible non-dominated set versus the number of evaluations: (a) optimisers on the co-design problem; (b) design strategies. Lines are means over runs and bands span the minimum and maximum

Table 9. Hypervolume of the feasible non-dominated sets obtained with 5120 evaluations per run (mean $\pm$ standard deviation over runs): HV with reference point 1.1 in normalised objectives, HVB0 with B0 as reference point (volume of improvements over B0 in all three objectives), feasible share of all evaluated designs and front size

Strategy / algorithm	HV	HVB0× 10^3	Runs	Feasible	Front size
Design strategies					
Co-design, NSGA-II (proposed)	1.0862±0.0016	3.01±0.08	3	73.7%	635
Sequential (geometry, then AP)	1.0647±0.0033	2.33±0.04	3	80.4%	703
Fixed autopilot (geometry only)	1.0565	2.20	1	96.3%	1442
Co-design, $\tau = 0.18$ fixed	0.7675	0.05	1	73.1%	679
B0 (single point)	0.3941	0	–	–	1
Optimisers on the co-design problem					
NSGA-II	1.0862±0.0016	3.01±0.08	3	73.7%	635
NSGA-III	1.0788±0.0079	3.30±0.11	3	75.0%	713
AGE-MOEA	1.0840±0.0048	3.15±0.11	3	73.0%	612
SMS-EMOA	1.0859±0.0078	3.24±0.13	3	71.0%	1037
LHS random search	0.9119±0.0230	0.51±0.27	3	15.6%	34

7.6. Comparison of Optimisers

On the co-design problem itself the choice of evolutionary algorithm matters much less than the choice of strategy (Table 9, Fig. 8a). The mean hypervolumes of NSGA-II (1.0862±0.0016), NSGA-

III (1.0788±0.0079), AGE-MOEA (1.0840±0.0048) and SMS-EMOA (1.0859±0.0078) lie within 0.7% of each other. Latin-hypercube random search with the same budget reaches only 0.9119±0.0230, and its B0-referenced hypervolume (0.51×10^{-3}) is 83% below that of NSGA-II, so evolutionary search is needed to exploit the narrow region of efficient and safe designs. One-sided Mann–Whitney tests of NSGA-II against the other optimisers give $p = 0.20$ (NSGA-III), $p = 0.50$ (AGE-MOEA), $p = 0.65$ (SMS-EMOA), $p = 0.05$ (LHS random search). NSGA-II reaches 95% of its final hypervolume after 1173 evaluations on average, i.e. well within the budget.

7.7. Ablation Study

Each ablation removes one modelling element during optimisation; the resulting front is then re-evaluated with the complete model (Table 10). Stall model. Optimising with a purely linear normal-force law gives a front that is unchanged when re-evaluated with the stall model (HV 1.0858 before and 1.0858 after re-evaluation, 100% of the designs feasible). For this ship the binding IMO criterion along the whole front is the initial turning ability at $10°$ helm, which never approaches stall; stall affects only the turning-circle indices of the smallest, high-aspect-ratio rudders, which are not binding. The stall model therefore acts as a safeguard that confirms the optimum rather than shifting it. Fluctuation term. Without the second-order fluctuation term the optimiser cannot see the energy cost of rudder activity. It selects weaker filtering (filter constants down to 10.7 s, against 20.3 s on the front of the complete model) and a median gain of 0.74, and its front underestimates the true rudder-related power by 12.9 kW on average and by up to 73.3 kW; re-evaluated with the complete model it reaches HV$= 1.0771$, below the 1.0858 of the complete-model optimisation with the same seed. Single design sea state. Designs optimised for the most probable sea state alone (cluster 1) are judged on milder conditions than the climate; re-evaluated over the full climate only 29.3% of them remain feasible, mainly because 70.7% exceed the steering-gear rate limit in the rougher sea states; their power rises by 40.1 kW on average and their heading deviation by 0.279°, and the hypervolume of the feasible part is 1.0180 against 1.0858 for the complete model. Adverse-weather constraint. Without the MEPC course-keeping requirement the optimiser shrinks the rudder down to 50% of the original area and claims a minimum power of 209.6 kW, but only 79.9% of these designs satisfy the requirement when it is checked, and the hypervolume of the feasible part is 1.0751 against 1.0858 for the complete model. The adverse-weather constraint is thus the element that keeps the energy optimisation from producing rudders that cannot hold course in heavy weather.

Table 10. Ablation study (all runs with the same seed and budget). Each variant is optimised with its simplified evaluation; its front is then re-evaluated with the complete model. Δf_E and Δf_ψ are the mean changes caused by the re-evaluation

Variant	Front size	HV (claimed)	HV (re-evaluated)	Feasible after re-evaluation	Δf_E (kW)	Δf_ψ (deg)
Complete model (proposed)	692	1.0858	1.0858	100%	–	–
No stall model	692	1.0858	1.0858	100.0%	+0.0	+0.000
No fluctuation term	1249	1.0702	1.0771	100.0%	+12.9	+0.000
Single design sea state	962	1.6493	1.0180	29.3%	+40.1	+0.279
No adverse-weather constraint	632	1.0912	1.0751	79.9%	+0.0	+0.000

7.8. Sensitivity Analysis

7.8.1. Global Sensitivity to the Design Variables

Figure 9 shows first-order and total Sobol indices over the full design box. The IMO utilisation is governed by the rudder area ($S_1 = 0.83$), followed by the section thickness ($S_1 = 0.12$) and the span ($S_1 = 0.04$); the autopilot has no influence because the IMO manoeuvres are open-loop. All course-keeping quantities are controlled by the autopilot: the filter time constant dominates the steering-gear

rate demand ($S_T = 0.73$) and the adverse-weather rudder demand ($S_T = 0.73$), and the gain dominates the heading deviation ($S_T = 0.58$). Over the full box the energy objective is also dominated by the controller ($S_T = 0.76$ for T_f, 0.64 for K_p and 0.40 for T_d), with total indices far above the first-order ones: the controller parameters act through strong interactions, because a weakly filtered, high-gain autopilot can multiply the steering losses. Geometry and controller therefore cannot be designed independently.

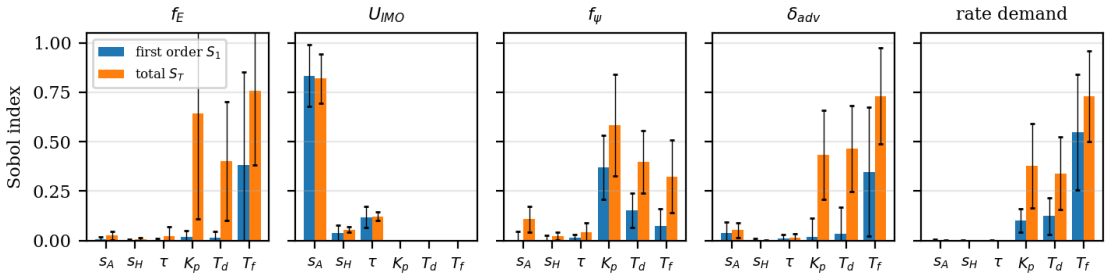


Figure 9. First-order and total Sobol indices of the objectives and constraint metrics with respect to the six design variables (Saltelli sampling, $N = 256$, 3584 evaluations; bars show 95% confidence intervals)

7.8.2. Sensitivity to Modelling Assumptions

Table 11 varies the stall parameters, the windage area, the steering-gear and integral time constants, the number of climate clusters and the wind–wave dependence one at a time. Under every variant the knee design dominates B0, and its saving stays between 14.9% and 15.1%. The stall parameters affect neither B0 nor the knee design, whose effective angles of attack (at most about $32°$ in the turning circle) remain below their stall angles; they matter only for the minimum-power design, whose IMO utilisation rises from 0.770 to 0.799 when the baseline stall angle is reduced to $33°$. As expected for designs on an active constraint, the designs at the low-energy end are the most sensitive: a 30% larger windage raises the adverse-weather demand of the minimum-power design beyond the limit. This motivates the robust formulation.

Table 11. Sensitivity of B0, the knee design and the minimum-power design to modelling assumptions (one at a time)

	f_E (kW)		Saving	U_{IMO}		Min-f_E design
Variant	B0	Knee	of knee	B0	Knee	U_{IMO} / δ_{adv}
Nominal model	315.2	267.8	15.0%	0.705	0.691	0.770 / 25.0°
$\alpha_{s0} = 33°$	315.2	267.8	15.0%	0.705	0.691	0.799 / 25.0°
$\alpha_{s0} = 41°$	315.2	267.8	15.0%	0.705	0.691	0.770 / 25.0°
$k_p = 0.10$	315.2	267.8	15.0%	0.705	0.691	0.770 / 25.0°
$k_p = 0.40$	315.2	267.8	15.0%	0.705	0.691	0.771 / 25.0°
windage $\times 0.7$	313.6	266.4	15.1%	0.705	0.691	0.770 / 24.4°
windage $\times 1.3$	316.8	269.4	15.0%	0.705	0.691	0.770 / 25.6°
$T_E = 1.5$ s	317.9	270.5	14.9%	0.705	0.691	0.770 / 24.9°
$T_E = 4.0$ s	314.7	267.6	15.0%	0.705	0.691	0.770 / 25.1°
$T_I = 200$ s	315.6	268.3	15.0%	0.705	0.691	0.770 / 25.2°
$T_I = 800$ s	315.0	267.6	15.0%	0.705	0.691	0.770 / 24.9°
climate, 6 clusters	312.3	265.2	15.1%	0.705	0.691	0.770 / 25.0°
climate, 12 clusters	317.1	269.7	15.0%	0.705	0.691	0.770 / 25.0°
climate, independent wind/waves	315.0	267.6	15.0%	0.705	0.691	0.770 / 25.0°
climate, copula $\rho = 0.5$	317.7	270.2	14.9%	0.705	0.691	0.770 / 25.0°

7.9. Robustness Under Hydrodynamic Uncertainty

The robust co-design of Eq. (22) was verified together with the nominal front and B0 by Monte Carlo simulation with 200 independent realisations of the coefficients of Table 2 (Fig. 10, Table 12). The adverse-weather rudder demand is the quantity that reacts most strongly to coefficient uncertainty: its 90% quantile for B0 is $31.9°$ against a nominal value of $21.1°$, so B0 meets both safety requirements in only 65.5% of the realisations. The designs of the nominal front reach a mean reliability of 0.769, with values as low as 0.435 at the low-energy end. The robust front raises the mean reliability to 0.910 (minimum 0.830), and 63% of its designs exceed 0.9 compared with 33% on the nominal front. Robustness does not come with a uniform energy penalty: for a reliability target of 0.9, the most efficient robust design requires 279.3 kW, against 290.6 kW for the most efficient nominal design that meets the same target and 315.2 kW for B0, which does not meet it.

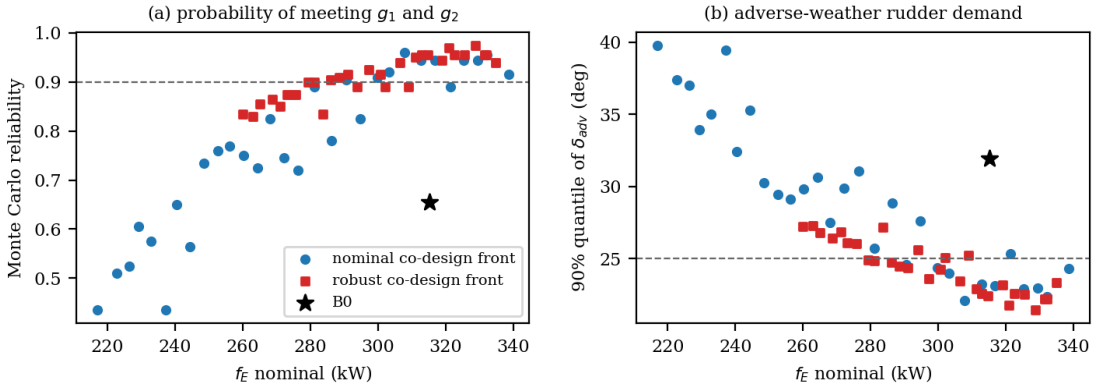


Figure 10. Monte Carlo verification (200 realisations of the hydrodynamic coefficients) of 30 designs each from the nominal and the robust co-design fronts and of B0: (a) probability of satisfying the IMO and adverse-weather requirements; (b) 90% quantile of the adverse-weather rudder demand

Table 12. Monte Carlo verification with 200 independent realisations of the hydrodynamic coefficients (30 designs per front). Reliability is the probability of satisfying g_1 and g_2

Design set	Mean reliability	Min. reliability	Share of designs with reliability ≥ 0.9	$Q_{0.9}(\delta_{adv})$ mean (deg)	Min. f_E with reliability ≥ 0.9
B0	0.655	0.655	0%	31.9	–
Nominal front	0.769	0.435	33%	29.0	290.6 kW
Robust front	0.910	0.830	63%	24.4	279.3 kW

8. Discussion

8.1. Where the Rudder Consumes Energy

For the reference design, the profile drag of the rudder in the propeller slipstream accounts for 265 kW of the 315 kW of rudder-related effective power (84%), the steady weather helm for 53 kW and the second-order fluctuation term for -3 kW. The near-zero fluctuation term is a consequence of the strong measurement filtering selected by the conventional tuning: the rudder then moves predominantly with the slowly varying yaw disturbance, the hull terms with negative and positive surge curvature largely cancel, and the cross-covariance between rudder angle and local cross-flow reduces the rudder normal force. This balance is fragile. For the same rudder, an autopilot with $K_p = 4$, $T_d = 80$ s and $T_f = 5$ s raises the fluctuation loss to 812 kW, 2.6 times the total rudder-related power of B0, because the derivative action then reacts to wave-frequency yaw rate. The Sobol analysis generalises this observation: over the whole design space the controller parameters carry the total-effect indices of the energy objective (Fig. 9), whereas on the Pareto front, where all controllers

are well filtered, the rudder area alone explains the energy level (correlation coefficient 0.999). The energy benefit of rudder geometry optimisation is therefore only realised if the autopilot is designed together with it; a geometry optimised under an aggressive controller would be judged on losses that a reasonable controller would never produce.

8.2. Design Rules for the Rudder of A Large Tanker

The Pareto set yields four rules. (i) Use the full available span. All non-dominated designs use a span within 1.2% of the upper bound; the higher aspect ratio raises the lift-gradient coefficient, so a smaller area delivers the same manoeuvring forces with less profile drag. The price is a lower stall angle ($34.5°$ for the knee design), which the stall-aware model shows to be harmless at moderate area; only very small, high-aspect-ratio rudders reach it, and for the minimum-power design (stall angle 28.9°) stall lengthens the advance of the $35°$ turning circle by $0.10L$ (3%) relative to a non-stalling rudder. (ii) Use the thinnest section the structure permits. The thin NACA section combines a higher normal-force slope with lower profile drag, and for a 15.5-kn tanker the cavitation-inception angle (about $9.5°$) remains well above the service rudder demand ($\delta_{95} \approx 3.0 - 4.7°$). The step-by-step decomposition from B0 to the knee design (Table 7) attributes -38.8 kW to the section, -6.9 kW to the span and -2.3 kW to the area. (iii) Size the area for adverse-weather course keeping. The IMO calm-water criteria are far from binding (utilisation 0.65–0.77 across the front); the low-energy end of the front is cut off by the rudder demand in the MEPC adverse conditions, which reaches its limit of $25°$ at the minimum-power design. This supports the view that manoeuvrability in adverse conditions must accompany energy-efficiency measures [3, 4]. (iv) Filter strongly and let the integral action carry the weather helm. All front designs use filter time constants of 20–40 s, which keep the rate demand within the capability of the steering gear; gain and derivative time then trade heading precision against rudder activity.

8.3. Methodological Implications

The spectral evaluator is the enabling element of the study. A time-domain assessment of one design over the 56 service conditions with records long enough to estimate the mean resistance within about 10% requires roughly 47 s of computation per random seed, so the 5120 evaluations of a single optimisation run would take about 67 h; with the spectral evaluator the complete run, including all IMO manoeuvres and the adverse-weather analysis, takes about four minutes. Because the method reproduces the ranking of designs exactly in the verification and the scatter of its resistance estimate equals the statistical uncertainty of the time-domain reference, the speed-up does not come at the price of decision quality. The baseline-anchored rudder model plays a complementary role: by construction it reproduces the validated behaviour of the original rudder, and its geometric extrapolation is governed by ratios of physically grounded quantities rather than by absolute empirical estimates.

8.4. Scope and Extensions

The study considers a single-screw tanker in deep water at service speed, a semi-balanced rudder parameterised by span, area and NACA thickness, and long-crested co-directional wind and waves. Natural extensions are twin-rudder and high-lift rudder configurations, rudder-stock and steering-gear torque requirements as additional structural constraints, wave-drift loads computed with panel methods, and ship types with larger windage, such as container ships and car carriers, for which the weather-helm share of the rudder-related power is expected to be larger. The framework itself is ship-independent: it requires an MMG coefficient set, section data and an operational climate.

9. Conclusions

This paper presented a spectral co-design method for ship rudder systems that optimises rudder planform, section and heading autopilot simultaneously for energy efficiency and manoeuvring safety.

Applied to the KVLCC2 tanker with captive-model data, free-running test data and NOAA buoy records, it leads to the following conclusions.

The baseline-anchored, stall-aware rudder model reproduces the reference MMG predictions of the original ship (mean differences $0.21L$ in the turning indices and $0.76°$ in the overshoot angles) and extrapolates consistently to new geometries; stall becomes relevant only for small, high-aspect-ratio rudders.

The spectral evaluator reproduces nonlinear time-domain statistics with $R^2 = 0.996$ for heading and rudder motions and $R^2 = 0.976$ for the steering resistance, preserves the ranking of designs exactly and is about 1100 times faster, which makes climate-based multi-objective co-design feasible on a single CPU core.

Against a reference design with the original rudder and a conventionally tuned autopilot, 187 Pareto-optimal designs improve all three objectives simultaneously. The knee design reduces the rudder-related power by 15.0% while lowering the IMO utilisation from 0.705 to 0.691 and the heading deviation from $0.255°$ to $0.240°$; the most energy-efficient design that still dominates the reference saves 18.9%.

Simultaneous co-design outperforms the conventional strategies at equal computational budget: its hypervolume exceeds that of sequential design by 2.0% and that of geometry optimisation with a fixed autopilot by 2.8%, and its B0-referenced hypervolume is 1.3 and 1.4 times larger. The choice among modern evolutionary algorithms is secondary (hypervolumes within 0.7%). The thinner section is the largest single lever: with the original NACA 0018 section retained, the most efficient design that dominates B0 saves 3.6%. Nonlinear time-domain simulation over the full climate confirms the savings (15.4% for the knee design).

The minimum rudder size is governed by course keeping in adverse conditions rather than by the IMO calm-water criteria; under hydrodynamic-coefficient uncertainty this requirement determines the reliability of a design. The quantile-based robust co-design raises the Monte Carlo reliability of the Pareto designs from 0.77 to 0.91 on average, and for a reliability target of 0.9 it requires 279.3 kW, 3.9% less than the best nominal design meeting the same target and 11.4% less than B0, whose reliability is only 0.66.

The method provides naval architects with a quantitative basis for sizing rudders and tuning autopilots together, and with an explicit picture of the manoeuvring margin that each kilowatt of saving costs.

Acknowledgements

The author acknowledges the Japan Society of Naval Architects and Ocean Engineers for the publication of the MMG standard method and the KVLCC2 captive-model data, MARIN and the SIMMAN organisers for the free-running model-test data, and the NOAA National Data Buoy Center for the buoy records.

Declarations

Funding

Not applicable.

Conflict of interest

The author declares no competing interests.

Data availability

The KVLCC2 hydrodynamic coefficients and free-running test results are published in [10, 13, 20]. The buoy records are available from the NOAA National Data Buoy Center [18]; the files used here

were obtained from the example-data repository of the MHKiT-Python toolkit. All derived data (section tables, climate clusters, optimisation archives and Monte Carlo results) are provided with the code.

Code availability

The complete Python implementation (MMG simulator, rudder model, spectral evaluator, optimisation drivers and analysis scripts) is provided as supplementary material.

Author contributions

Fuyong Zhou is the sole author and carried out the conceptualisation, methodology, software development, analysis and writing.

KVLCC2 Model Parameters

Table A1 lists the parameters of the MMG model used in this study. Hydrodynamic coefficients are non-dimensionalised with $1/2\,\rho L_{pp} dU^2$ (forces) and $1/2\,\rho L_{pp}^2 dU^2$ (moment).

Table A1. MMG parameters of the KVLCC2 from captive-model tests [13] and propeller open-water polynomial [20]

R'_0	0.022	Y'_v	-0.315	N'_v	-0.137
X'_{vv}	-0.040	Y'_r	0.083	N'_r	-0.049
X'_{vr}	0.002	Y'_{vvv}	-1.607	N'_{vvv}	-0.030
X'_{rr}	0.011	Y'_{vvr}	0.379	N'_{vvr}	-0.294
X'_{vvvv}	0.771	Y'_{vrr}	-0.391	N'_{vrr}	0.055
m'_x	0.022	Y'_{rrr}	0.008	N'_{rrr}	-0.013
m'_y	0.223	J'_z	0.011	x'_P	-0.48
t_P	0.220	w_{P0}	0.40	C_1	2.0
C_2 $(\beta_P > 0/< 0)$	1.6/1.1	t_R	0.387	a_H	0.312
x'_H	-0.464	x'_R	-0.500	γ_R $(\beta_R < 0/> 0)$	0.395/0.640
l'_R	-0.710	ε	1.09	κ	0.50
$f_{\alpha 0}$	2.747	Λ_0	1.827	K_T	$0.2931 - 0.2753J - 0.1385J^2$

References

[1] International Maritime Organization. (2002). Resolution MSC.137(76): Standards for ship manoeuvrability.

[2] International Maritime Organization. (2021). MEPC.1/Circ.850/Rev.3: Guidelines for determining minimum propulsion power to maintain the manoeuvrability of ships in adverse conditions.

[3] Papanikolaou, A., Zaraphonitis, G., Bitner-Gregersen, E., et al. (2016). Energy efficient safe ship operation (SHOPERA). Transportation Research Procedia, 14, 820–829.

[4] Shigunov, V., & Papanikolaou, A. (2015). Criteria for minimum powering and maneuverability in adverse weather conditions. Ship Technology Research, 62, 140–147. https://doi.org/10.1080/09377255.2015.1104090

[5] Sutulo, S., & Guedes Soares, C. (2021). Review on ship manoeuvrability criteria and standards. Journal of Marine Science and Engineering, 9, 904. https://doi.org/10.3390/jmse9080904

[6] Molland, A. F., & Turnock, S. R. (2007). Marine rudders and control surfaces: Principles, data, design and applications. Butterworth-Heinemann.

[7] Liu, J., & Hekkenberg, R. (2017). Sixty years of research on ship rudders: Effects of design choices on rudder performance. Ships and Offshore Structures, 12, 495–512.

[8] Koyama, T. (1967). On the optimum automatic steering system of ships at sea (in Japanese). Journal of the Society of Naval Architects of Japan, 1967.

[9] Reid, R. E. (1978). Design of an automatic steering control system for ships to minimize added resistance due to steering [Doctoral dissertation]. University of Virginia.

[10] Liu, J., Quadvlieg, F., & Hekkenberg, R. (2016). Impacts of the rudder profile on manoeuvring performance of ships. Ocean Engineering, 124, 226–240. https://doi.org/10.1016/j.oceaneng.2016.07.064

[11] Fossen, T. I. (2021). Handbook of marine craft hydrodynamics and motion control (2nd ed.). John Wiley & Sons.

[12] Fujii, H., & Tsuda, T. (1961). Experimental researches on rudder performance (2) (in Japanese). Journal of Zosen Kiokai, 110, 31–42.

[13] Yasukawa, H., & Yoshimura, Y. (2015). Introduction of MMG standard method for ship maneuvering predictions. Journal of Marine Science and Technology, 20, 37–52. https://doi.org/10.1007/s00773-014-0293-y

[14] Stern, F., Agdrup, K., Kim, S. Y., et al. (2011). Experience from SIMMAN 2008—the first workshop on verification and validation of ship maneuvering simulation methods. Journal of Ship Research, 55, 135–147. https://doi.org/10.5957/jsr.2011.55.2.135

[15] Sharpe, P., & Hansman, R. J. (2025). https://doi.org/10.48550/arXiv.2503.16323

[16] Yasukawa, H., & Sakuno, R. (2020). Application of the MMG method for the prediction of steady sailing condition and course stability of a ship under external disturbances. Journal of Marine Science and Technology, 25, 196–220. https://doi.org/10.1007/s00773-019-00641-4

[17] Newman, J. N. (1974). Second-order, slowly-varying forces on vessels in irregular waves. In Proceedings of the International Symposium on Dynamics of Marine Vehicles and Structures in Waves (pp. 182–186).

[18] National Data Buoy Center. (2024). Historical standard meteorological, continuous wind and wave data for stations 46002, 46050 and 46097.

[19] Quadvlieg, F. H. H. A., & Brouwer, J. (2011). KVLCC2 benchmark data including uncertainty analysis to support manoeuvring predictions. In Proceedings of the 4th International Conference on Computational Methods in Marine Engineering (MARINE 2011).

[20] Dai, K., & Li, Y. (2019). Manoeuvring prediction of KVLCC2 with hydrodynamic derivatives generated by a virtual captive model test. Polish Maritime Research, 26, 16–26. https://doi.org/10.2478/pomr-2019-0062

[21] Faltinsen, O. M., Minsaas, K. J., Liapis, N., & Skjørdal, S. O. (1980). Prediction of resistance and propulsion of a ship in a seaway. In Proceedings of the 13th Symposium on Naval Hydrodynamics (pp. 505–529).

[22] Fujii, H., & Takahashi, T. (1975). Experimental study on the resistance increase of a ship in regular oblique waves. In Proceedings of the 14th International Towing Tank Conference.

[23] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6, 182–197.

[24] Deb, K., & Jain, H. (2014). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. IEEE Transactions on Evolutionary Computation, 18, 577–601.

[25] Beume, N., Naujoks, B., & Emmerich, M. (2007). SMS-EMOA: Multiobjective selection based on dominated hypervolume. European Journal of Operational Research, 181, 1653–1669.

[26] Panichella, A. (2019). https://doi.org/10.1145/3321707.3321839. In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '19) (pp. 595–603).

[27] Blank, J., & Deb, K. (2020). Pymoo: Multi-objective optimization in Python. IEEE Access, 8, 89497–89509.

[28] Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. IEEE Transactions on Evolutionary Computation, 3, 257–271.

[29] Beyer, H.-G., & Sendhoff, B. (2007). Robust optimization – a comprehensive survey. Computer Methods in Applied Mechanics and Engineering, 196, 3190–3218.

[30] Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation, 55, 271–280.

[31] Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. Computer Physics Communications, 145, 280–297.

[32] Blendermann, W. (1994). Parameter identification of wind loads on ships. Journal of Wind Engineering and Industrial Aerodynamics, 51, 339–351.

[33] Det Norske Veritas. (2010). Recommended practice DNV-RP-C205: Environmental conditions and environmental loads.

[34] Hasselmann, K., Barnett, T. P., Bouws, E., et al. (1973). Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). Ergänzungsheft zur Deutschen Hydrographischen Zeitschrift, Reihe A.

[35] Herman, J., & Usher, W. (2017). SALib: An open-source Python library for sensitivity analysis. Journal of Open Source Software, 2, 97.