UCLouvain

epl

École polytechnique de Louvain

Hull Optimization for Fast and Efficient Semi-Planing Craft using RANS CFD combined with the Volume of Fluid Method

Author: **Gauthier GHEYSENS**
Supervisor: **Grégoire WINCKELMANS**
Readers: **Denis-Gabriel CAPRACE, Yann BARTOSIEWICZ**
Academic year 2025–2026
Master [120] in Mechanical Engineering

Abstract

This thesis presents a fully automated computational fluid dynamics (CFD) optimization pipeline developed for the hull form of a fast semi-planing electric yacht (Lumen Yachts, $L_{wl} = 7\,\mathrm{m}$, $Fr = 0.796$, corresponding to twice the hull speed). The pipeline integrates parametric hull generation through Grasshopper and Rhino.Compute, fluid domain construction in FreeCAD and SpaceClaim, mesh generation and flow simulation in ANSYS Meshing and Fluent, and result extraction via Python scripts. It constitutes an optimization tool: any geometric parameter can be targeted and the entire evaluation chain runs without manual intervention. The hull is a 7-metre chined semi-planing craft, operating in a regime where both wave-making and dynamic lift are significant.

The flow is modelled using Reynolds-Averaged Navier-Stokes (RANS) equations with the realizable k-ε turbulence model and scalable wall functions, combined with the Volume of Fluid (VOF) method for free-surface treatment. A mesh sensitivity study confirms that the 2.7% drag deviation between the selected and finest meshes is small compared to the drag variations between hull variants, placing the selected mesh in a suitable operating range for comparative hull evaluation.

A one-dimensional Sobol exploration over the stern chine position, the parameter identified as the most influential on drag, reveals a flat optimal zone Chine_Y $\in$ [0.058, 0.080] where the mean drag is 6.7% below the base design. Within this zone, drag differences between hull variants fall below the numerical noise floor, so no single point can be identified as the unique optimum. The nominal minimum (Chine_Y $= 0.071$, $D = 1359.8\,\mathrm{N}$, -7.7% relative to base) is within this uncertainty and should be understood as one representative point in the optimal zone rather than an individually identified optimum. The optimization infrastructure supports multi-parameter Bayesian optimization, demonstrated on preliminary runs, but the Sobol phase alone was sufficient to characterize the 1D drag landscape.

The results confirm that hull geometry at the chine has a strong and non-trivial influence on resistance in the semi-planing regime. The automated framework developed here provides a reusable foundation for future optimization campaigns with larger computational budgets.

Keywords: Semi-planing hull, RANS CFD, Bayesian optimization, hull form optimization, Volume of Fluid, wall functions, mesh convergence, Lumen Yachts

Acknowledgements

I would like to express my sincere gratitude to Professor Grégoire Winckelmans, my thesis supervisor, for his guidance, his availability, and the depth of knowledge in fluid mechanics he shared throughout this work. His remarks were always precise and constructive, and his expertise in CFD shaped this project.

I am equally grateful to Jean-Baptiste Crismer, who helped me get to grips with the computational infrastructure and offered advice on how to structure and present the work. His involvement went well beyond what could reasonably be expected.

I also thank the readers, Professor Denis-Gabriel Caprace and Professor Yann Bartosiewicz, for agreeing to evaluate this work and for the time they devoted to it.

I would also like to thank Lumen Yachts for the opportunity to work on this project. Designing a fast electric yacht with a reduced environmental footprint is a genuinely meaningful engineering challenge, and being able to contribute to it through this thesis made the work all the more rewarding.

Finally, I would like to thank my family for their constant support throughout my studies. Their encouragement and patience, particularly during the more demanding phases of this thesis, made a real difference.

Contents

1 Introduction **11**

1.1 Context and Motivation 11

1.2 Objectives and Contribution 11

1.3 Literature Survey 12

1.3.1 CFD-Based Hull Optimization 12

1.3.2 Semi-Planing Hull Hydrodynamics 12

1.3.3 Automation and Parametric Geometry 13

1.4 Thesis Structure 13

2 Theoretical Background **14**

2.1 Ship Resistance Fundamentals 14

2.1.1 Resistance Components 14

2.1.2 The Semi-Planing Transition 14

2.2 RANS Equations 15

2.2.1 Reynolds Decomposition and Averaging 15

2.2.2 Boussinesq Hypothesis 16

2.3 Turbulence Models 17

2.3.1 The Standard and Realizable k-ε Models 17

2.3.2 The k-ω SST Model 17

2.4 Near-Wall Treatment and Wall Functions 18

2.4.1 Structure of the Near-Wall Region 18

2.4.2 Standard Wall Functions 18

2.4.3 Scalable Wall Functions 19

2.4.4 Wall Function Limitations in Adverse Pressure Gradients . . . 19

2.4.5 Model and Near-Wall Treatment Selection 20

2.5 Scale Effects and Full-Scale Strategy 21

2.6 Free-Surface Modelling: VOF Method 21

3 Numerical Setup **23**

3.1 Hull Geometry and Parameterization 23

3.1.1 Vessel Description 23

3.1.2 Parametric Hull Generation with Grasshopper 24
3.2 Computational Domain 25
3.2.1 Boundary Conditions 25
3.3 Mesh Generation 26
3.4 Solver Settings 27
3.5 Hull Position: Trim and Heave Equilibrium 28
3.6 Convergence Monitoring for Individual CFD Evaluations 29

4 Automated Optimization Framework **31**
4.1 Parameter Sensitivity Analysis 31
4.1.1 Motivation 31
4.1.2 Pearson Correlation Analysis 31
4.1.3 Choice of Optimization Parameter 33
4.2 Overview of the Pipeline 33
4.3 Trim and Heave Equilibrium Algorithm 34
4.3.1 Equilibrium Solver: Newton-Raphson with Finite-Difference Jacobian 34
4.3.2 Warm Start Regression Model 35
4.4 Optimization Algorithms 36
4.4.1 Bayesian Optimization Principles 36
4.4.2 Initialization: Sobol Sequences 36
4.4.3 Multi-Parameter Bayesian Optimizer 37
4.5 Development Path and Abandoned Approaches 37
4.5.1 GMSH for Volume Mesh Generation 37
4.5.2 Transient Simulation with Free Hull Motion 38

5 Verification and Numerical Assessment **39**
5.1 Mesh Sensitivity Study 39
5.1.1 Context and Methodology 39
5.1.2 Mesh Configurations 40
5.1.3 Results and Selected Mesh 41
5.2 Boundary Layer Quality 41
5.2.1 y^+ Distribution 41
5.2.2 Boundary Layer Profile: Comparison with the Law of the Wall 42
5.3 Numerical Accuracy and Sensitivity Assessment 44
5.3.1 Two Distinct Sources of Error: Discretization Imprecision and Model Bias 44
5.3.2 Sensitivity to Boundary Conditions 45
5.4 Comparison with Reference Literature 47
5.4.1 Order-of-Magnitude Comparison with Molland et al. (2011) . 47

6 Results and Discussion **48**
6.1 Optimization Campaign . . . 48
6.1.1 Evaluations Performed . . . 48
6.1.2 Drag as a Function of the Chine Parameter . . . 50
6.2 Trim Equilibrium Summary . . . 51
6.2.1 Statistics Across All Evaluations . . . 51
6.2.2 Example Convergence: Evaluation 7 . . . 52
6.3 Flow Field Analysis at the Optimal Configuration . . . 53
6.4 Best Hull Geometry and Comparison with the Base Design . . . 55
6.4.1 Free-Surface Elevation . . . 55
6.4.2 Wall Shear Stress . . . 57
6.4.3 Pressure Distribution on the Hull Surface . . . 57
6.4.4 Comparison with the Original Lumen Yachts Hull . . . 58
6.5 Discussion . . . 59
6.5.1 Physical Interpretation . . . 59
6.5.2 Comparison to Bretscher (2017) . . . 60
6.5.3 Limitations and Sources of Uncertainty . . . 60

7 Conclusion **61**
7.1 Summary and Scientific Perspective . . . 61
7.2 Limitations . . . 62
7.3 Perspectives . . . 62

A Nomenclature and Constants **67**

B Optimization Results Summary **68**

C Trim Equilibrium Convergence — All Evaluations **69**

D CFD Solver Convergence Histories **72**

E Full Mesh Sensitivity Table **73**

F Grasshopper Hull Parameters **74**

G Velocity Streamlines — Base Hull (Evaluation 4) **76**

H Declaration of Use of Artificial Intelligence Tools **77**

List of Figures

2.1 The three speed regimes for a planing hull as a function of Froude number: displacement (left), semi-planing/semi-displacement (centre), and planing (right). Source: Molland et al. [18]. 15

3.1 Grasshopper parametric model of the Lumen Yachts hull. Left: a portion of the Grasshopper canvas showing a selection of the parameter sliders. Right: the resulting hull geometry displayed in the Rhinoceros viewport. The parametric definition exposes 29 geometric variables controlling the hull shape; `Chine_stern_point_Y` (the lateral chine position at the stern) is the optimization parameter. The full parameter list is given in Appendix F. 24

3.2 Lateral cross-section of the computational domain. The body of influence (BOI) is visible as a refined zone around the free surface. The hull is at the right; the inlet is at the far right, the outlet at the left. 26

3.3 Zoomed view of the boundary layer mesh near the hull surface. The 16 structured layers are clearly visible, with the first-cell height of $\Delta y_1 = 1\,\text{mm}$ at the wall, transitioning to polyhedral cells in the far field. 27

3.4 Forces and moments acting on the hull at speed. ANSYS Fluent computes the hydrodynamic drag $F_{Trainée}$ (longitudinal resistance), the vertical hydrodynamic forces $F_{Archimede}$ and $F_{Portance}$ (lift plus buoyancy), and the pitching moment M_y about the centre of gravity. The motor thrust $F_{Poussée}$ applied at the propeller shaft introduces an additional pitching moment $M_{\text{motor}} = F_{Poussée} \cdot d_z$, where d_z is the vertical offset between the thrust line and the centre of gravity. 29

4.1 Pearson correlation coefficients for the 29 hull parameters. The top 10 are labelled. Chine_stern_point_Y (p8) is clearly dominant. . . . 32

4.2 Automated optimization pipeline. Each block is one software tool; all transitions are driven by Python scripts. The red dashed arrow shows the trim/heave equilibrium iteration: if force and moment residuals exceed the threshold, the hull position is updated and a new CFD simulation is launched. The green arrow returns the converged drag to the optimizer. 33

5.1 Percentage deviation of drag coefficient C_D relative to the finest mesh (1.5 cm reference) as a function of hull face size Δx. The orange marker indicates the selected mesh ($\Delta x = 3$ cm, +2.71% deviation). Richardson extrapolation is not applied: scalable wall functions impose a fixed physical length scale at the wall that does not reduce with bulk mesh refinement, making the observed convergence order unreliable for extrapolation (see Section 5.3). 42

5.2 y^+ distribution on the hull surface for the selected mesh (3 cm face size). Values on the wetted hull surface reach a maximum of approximately 150; low values ($y^+ \approx 5$) appear in the air domain above the waterline. All submerged-surface values fall within the log-layer range $30 < y^+ < 300$, confirming correct scalable wall function operation. 43

5.3 Boundary layer velocity profile extracted from CFD (symbols) compared with the log law using modern constants $\kappa = 0.383$, $B = 4.25$. Four points follow the log law; the three outer points depart upward into the wake region. 44

6.1 Drag force (top), equilibrium trim angle (middle), and equilibrium heave (bottom) as functions of Chine_stern_point_Y. All 10 evaluations converged to the 50 N tolerance (circles). Each point is labelled with its evaluation index (e1–e11 in chronological order; e8 is absent: it was excluded for the reasons detailed in Section 6.2). The best hull (Chine_Y = 0.071, $D = 1359.8$ N, evaluation 7) is shown in orange; the base design (Chine_Y = 0.106, $D = 1472.5$ N, evaluation 4) in red. The green dashed box marks the ideal zone (Chine_Y ≤ 0.080, 6 evaluations), with a mean drag of 1373.7 N, which is 6.7% below the base design. The same zone is highlighted on the trim and heave subplots. 49

6.2 Final trim equilibrium residuals for all 10 evaluations, sorted by Chine_Y ascending. Blue bars: force residual $|F_z - W|$ (N). Red bars: pitching moment residual $|M_y|$ (N·m). The dashed line marks the 50 N / 50 N·m convergence threshold. The number of trim iterations n is indicated below each bar group. All evaluations satisfy both criteria simultaneously. 52
6.3 Trim equilibrium convergence for evaluation 7 (Chine_Y = 0.071, best hull). Top row: force residual $|F_z - W|$ (left) and pitching moment residual $|M_y|$ (right); the green band marks the convergence zone below the 50 N / 50 N·m threshold (dashed line). Bottom row: evolution of the trim angle α (left) and heave z (right) at each Newton-Raphson iteration, converging to $\alpha = 5.89°$ and $z = 156.6$ mm by iteration 14. . 53
6.4 VOF free-surface rendering of the optimal hull (eval 7, Chine_Y = 0.071) at $Fr = 0.796$, quarter view from starboard. The bow wave, diverging wake, and clean water detachment at the transom are clearly visible. 54
6.5 Front view of the optimal hull at $Fr = 0.796$. The bow wave pushes water outward on both sides. 55
6.6 Free-surface Z elevation at $Fr = 0.796$ — optimal hull (eval 7, Chine_Y = 0.071): range −0.273 to +0.383 m. 56
6.7 Free-surface Z elevation at $Fr = 0.796$ — base hull (eval 4, Chine_Y = 0.106): range −0.225 to +0.391 m. 56
6.8 Wall shear stress τ_w [Pa] on the hull surface. Left: optimal configuration. Right: base configuration. The high-shear region is concentrated at the bow shoulder in both cases. 57
6.9 Static pressure [Pa] on the hull surface. Left: optimal configuration. Right: base configuration. 58

C.1 Trim equilibrium convergence — evaluations 11, 9, 10, 6, 7 (Chine_Y = 0.058 to 0.071). Left column: force residual $|F_z - W|$ (N). Right column: pitching moment residual $|M_y|$ (N·m). Dashed line: 50 N / 50 N·m threshold. Green band: convergence zone. 70
C.2 Trim equilibrium convergence — evaluations 1, 2, 3, 4, 5 (Chine_Y = 0.076 to 0.116). Same layout as Figure C.1. 71

D.1 CFD drag convergence history for the last trim iteration of each of the 10 optimization evaluations (eval 8 excluded). Each panel shows one hull (Chine_Y value indicated); the dashed line marks the recorded converged drag (Table 6.1). The colour scale maps Chine_Y from narrow chine (dark, purple) to wide chine (light, yellow). Note: for evaluations 5, 6, and 9, which were re-run to the 50 N criterion on a separate workstation, the figure shows the last locally available trim iteration; the converged drag values in Table 6.1 were taken from the completed re-run. 72

G.1 Velocity streamlines on the free surface for the base hull (evaluation 4, Chine_Y = 0.106). Colour scale: velocity magnitude in $\mathrm{m\,s^{-1}}$. 76

List of Tables

2.1 Summary comparison of turbulence models for marine CFD. 18

3.1 Key vessel parameters used in all simulations. 23

4.1 Top 10 Grasshopper parameters ranked by absolute Pearson correlation with drag (50 preliminary simulations at fixed trim/heave). Physical interpretations are given for the most influential parameters. 32

5.1 Mesh sensitivity study: percentage deviation of dimensionless coefficients relative to the finest mesh (1.5 cm reference). The selected mesh ($\Delta x = 3\,\text{cm}$) is highlighted. Base hull at fixed trim ($\alpha = 0°$). . . 40

6.1 All 10 optimization evaluations, sorted by Chine_Y (ascending). The best result (lowest drag) is highlighted in bold; the base design (Chine_Y = 0.106) is in italics. All evaluations converged to the 50 N / 50 N·m tolerance. n_{CFD} is the number of steady-state CFD simulations run for that evaluation. The evaluation index is the chronological order of execution, not a rank. 50

6.2 Trim equilibrium convergence statistics for all 10 evaluations. 51

A.1 Symbols and physical constants used throughout this thesis. 67

B.1 Complete optimization results in chronological order (evaluation index). All values correspond to the 50 N convergence criterion. Evaluation 8 is excluded (terminated prematurely, not representative of equilibrium). 68

E.1 Full mesh sensitivity table: absolute dimensionless coefficients and percentage deviation from the finest mesh (1.5 cm reference). The selected mesh ($\Delta x = 3\,\text{cm}$) is highlighted. Base hull at fixed trim ($\alpha = 0°$). 73

F.1 Grasshopper hull parameters (p_0–p_{28}). The optimised parameter (p_8) is highlighted in bold. 75

Chapter 1

Introduction

1.1 Context and Motivation

The maritime sector faces growing pressure to reduce fuel consumption and emissions. For electric vessels in particular, range anxiety makes hydrodynamic efficiency not merely desirable but fundamental: every newton of unnecessary drag translates directly into battery weight, reduced autonomy, or both. Lumen Yachts, a Dutch manufacturer of electric pleasure craft, is studying the development of a 7 m semi-planing hull intended to cruise at twice hull speed. At this operating point, the hull transitions from a displacement regime, where wave-making dominates, into a semi-planing regime where both wave-making and dynamic lift play significant roles. Optimizing hull geometry in this regime is therefore a meaningful engineering challenge.

Traditional hull design relies on towing tank testing and designer experience. While valuable, these approaches are expensive, time-consuming, and offer limited insight into the detailed flow physics driving resistance. Computational fluid dynamics, specifically RANS-based solvers, offers a cost-effective alternative that enables systematic parametric exploration of the design space. Recent advances in automation and surrogate-based optimization have made it realistic to couple CFD directly with geometry generation tools and optimization algorithms, reducing the manual effort per evaluation.

1.2 Objectives and Contribution

The contribution of this work is a fully automated CFD-based hull evaluation pipeline, applied to a chined semi-planing hull at $Fr = 0.796$, a semi-planing regime where both wave-making resistance and dynamic lift play significant roles.

The specific contributions are:

1. A fully automated end-to-end optimization pipeline integrating Grasshopper, Rhino.Compute, FreeCAD, SpaceClaim, ANSYS Meshing and Fluent, and Python-based optimization and post-processing codes.
2. A verified CFD setup (mesh convergence study, boundary layer quality assessment) for a semi-planing hull at $Fr = 0.796$.
3. A completed one-dimensional exploration campaign over the most influential geometric parameter, together with a multi-parameter Bayesian optimizer ready for deployment with greater computational resources.

1.3 Literature Survey

1.3.1 CFD-Based Hull Optimization

The coupling of RANS CFD with parametric hull geometry and optimization algorithms is an active field in naval architecture. Early work by Peri and Campana [19] demonstrated gradient-based multi-disciplinary optimization of a naval combatant using panel methods and RANS CFD as evaluation tools. Since then, the computational cost per evaluation has decreased substantially with hardware improvements, making surrogate-based methods increasingly practical.

Bretscher [5] demonstrated drag reductions through multi-parameter RANS-based optimization of a round-bilge pure-displacement hull. The present work applies a similar philosophy to a chined semi-planing hull at a substantially higher Froude number and with a fully automated pipeline.

Bayesian optimization with Gaussian Process surrogates has emerged as the preferred approach for expensive black-box optimization in engineering design, reviewed comprehensively by Shahriari et al. [24]. Its application to hull form optimization is growing but remains limited by the high cost of CFD evaluations.

1.3.2 Semi-Planing Hull Hydrodynamics

The transition from displacement to planing occurs roughly between $Fr \approx 0.4$ and $Fr \approx 1.0$ [18]. In the displacement regime ($Fr < 0.4$), wave-making resistance dominates and the hull is fully immersed. Above $Fr \approx 1.0$, the hull planes on its bottom with hydrodynamic lift carrying most of the weight, and friction on the wetted bottom becomes the primary resistance component. The semi-planing (or semi-displacement) regime between these two limits combines both mechanisms: dynamic lift is significant but buoyancy still plays a role, the transom may be partially ventilated, and both wave-making and friction contribute substantially to total resistance.

Le Strat and Terziev [16] conducted a detailed RANS study of the effects of trim on resistance for a sailing planing hull, demonstrating that equilibrium trim has a first-order influence on drag and that fixed-trim simulations systematically misrank hull variants. Their findings directly motivate the trim equilibrium algorithm developed in the present work.

1.3.3 Automation and Parametric Geometry

The use of Grasshopper/Rhinoceros as a parametric hull generation platform, driven programmatically via Rhino.Compute, represents a relatively recent development enabled by McNeel's HTTP API [23]. Combined with FreeCAD for domain construction and Python-driven ANSYS journal files, this provides a fully scriptable pipeline that eliminates the manual CAD steps that previously constituted the main bottleneck in optimization loops.

1.4 Thesis Structure

Chapter 2 establishes the theoretical background: RANS equations, turbulence closure, wall treatment, and scale effects. Chapter 3 describes the numerical setup, including geometry parameterization, domain, mesh, and solver settings. Chapter 4 details the automated optimization framework, from parameter sensitivity analysis to the Bayesian optimization algorithm. Chapter 5 presents the numerical verification and assessment of the CFD setup. Chapter 6 discusses the optimization results and their physical interpretation. Chapter 7 concludes with a summary, limitations, and directions for future work.

Chapter 2

Theoretical Background

2.1 Ship Resistance Fundamentals

2.1.1 Resistance Components

The total calm-water resistance R_T of a hull advancing at speed V is classically decomposed into two main components [18]:

$$R_T = R_V + R_W, \tag{2.1}$$

where R_V is the viscous resistance and R_W is the wave-making resistance.

Viscous resistance R_V arises from the shear stress exerted by the fluid on the hull surface. It is dominant at low Froude numbers and scales as $R_V \sim \rho V^2 S\, C_F(Re)$, where S is the wetted surface area and C_F is the friction coefficient.

Wave resistance R_W arises from the energy radiated in the form of surface waves. A hull advancing at constant speed generates the Kelvin wave system (transverse and divergent waves), which is the dominant wave pattern in displacement mode. The wave resistance is a strongly nonlinear function of Froude number: it rises steeply as the vessel approaches hull speed ($Fr \approx 0.4$), peaks near $Fr \approx 0.45$–0.50, then decreases as the hull enters the semi-planing and planing regimes and the wetted area reduces.

2.1.2 The Semi-Planing Transition

For a vessel operating at $Fr = 0.796$ (twice hull speed), the flow regime differs qualitatively from displacement mode. Three speed regimes are commonly identified: displacement ($Fr \lesssim 0.4$), semi-planing or semi-displacement ($0.4 \lesssim Fr \lesssim 1.0$), and fully planing ($Fr \gtrsim 1.0$) [10, 18]. The present hull operates in the semi-planing regime.

Three effects combine:

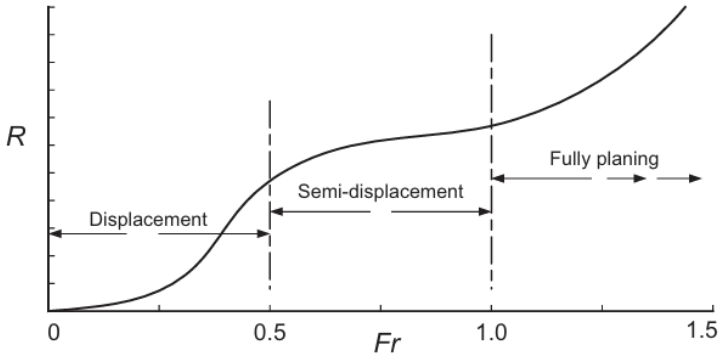


Figure 2.1: The three speed regimes for a planing hull as a function of Froude number: displacement (left), semi-planing/semi-displacement (centre), and planing (right). Source: Molland et al. [18].

Dynamic lift. The hydrodynamic pressure on the bottom of the hull generates a net upward force (lift) that supplements the static buoyancy. For the present hull at equilibrium, the CFD results show that a significant fraction of the weight is supported by dynamic lift: the hull trims bow-up ($\alpha \approx 6°$) and rises out of the water (heave $z \approx 147$ mm at the optimal configuration). This reduces the displaced volume and hence the wave resistance.

Wetted area reduction. As the hull rises, the wetted surface area decreases. This reduces both friction resistance and the effective form drag. However, the remaining wetted bottom, now acting as a planing surface, generates wave disturbances that contribute to wave resistance.

Spray resistance. At high speed, water is deflected laterally and upward at the chine, generating spray. The momentum transferred to this spray represents an additional resistance component not captured by classical wave-resistance theories. RANS-VOF simulations allow the large-scale free-surface deformation and spray structures to emerge naturally from the solution, but fine spray dynamics are only partially resolved due to mesh resolution and turbulence modelling limitations.

The total resistance at $Fr = 0.796$ is thus the result of a complex interaction between these three components, and the hull geometry at the stern chine (the optimized parameter) directly influences all three.

2.2 RANS Equations

2.2.1 Reynolds Decomposition and Averaging

Turbulent flows are governed by the Navier-Stokes equations. The Reynolds-Averaged Navier-Stokes (RANS) approach decomposes each flow variable into a mean and a

fluctuating component. For velocity:

$$u_i = \overline{u}_i + u_i', \qquad \overline{u_i'} = 0, \tag{2.2}$$

where $\overline{u}_i$ denotes the ensemble (Reynolds) average and u_i' the fluctuation. Substituting into the governing equations and averaging yields the RANS system. Following the notation of the ANSYS Fluent Theory Guide [2], the averaged continuity and momentum equations are:

$$\frac{\partial \rho}{\partial t} + \frac{\partial (\rho \overline{u}_i)}{\partial x_i} = 0, \tag{2.3}$$

$$\frac{\partial (\rho \overline{u}_i)}{\partial t} + \frac{\partial (\rho \overline{u}_i\, \overline{u}_j)}{\partial x_j} = -\frac{\partial \overline{p}}{\partial x_i} + \frac{\partial}{\partial x_j}\left[\mu\left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i} - \frac{2}{3}\delta_{ij}\frac{\partial \overline{u}_l}{\partial x_l}\right)\right] + \frac{\partial}{\partial x_j}\left(-\rho\,\overline{u_i' u_j'}\right), \tag{2.4}$$

where ρ is the fluid density, $\overline{p}$ is the mean pressure, μ is the dynamic viscosity, and $-\rho\,\overline{u_i' u_j'}$ is the *Reynolds stress tensor*. The viscous stress tensor includes the full symmetric form with the $2/3\,\delta_{ij}$ compressibility correction; for the incompressible flow considered here, $\partial \overline{u}_l / \partial x_l = 0$ and this term vanishes. The Reynolds stress tensor has six independent components not determined by the mean flow: this is the *closure problem* of turbulence [20]. Turbulence models provide the additional equations needed to close the system.

2.2.2 Boussinesq Hypothesis

The most widely used closure strategy relates the Reynolds stress tensor to the mean strain rate through a scalar eddy viscosity μ_t [1]:

$$-\rho\,\overline{u_i' u_j'} = \mu_t\left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i}\right) - \frac{2}{3}\left(\rho k + \mu_t \frac{\partial \overline{u}_k}{\partial x_k}\right)\delta_{ij}, \tag{2.5}$$

where $k = \frac{1}{2}\,\overline{u_i' u_i'}$ is the turbulent kinetic energy. For incompressible flow, $\partial \overline{u}_k / \partial x_k = 0$ and the last term reduces to $-\frac{2}{3}\rho k\,\delta_{ij}$. Substituting into Equation (2.4) yields a momentum equation formally identical to the laminar form with effective viscosity $\mu_{\text{eff}} = \mu + \mu_t$: the turbulence modelling problem reduces to providing a field μ_t [28].

2.3 Turbulence Models

2.3.1 The Standard and Realizable k-ε Models

The standard k-ε model [3,15] solves transport equations for turbulent kinetic energy k and its dissipation rate ε:

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \overline{u}_j k)}{\partial x_j} = \frac{\partial}{\partial x_j}\left[\left(\mu + \frac{\mu_t}{\sigma_k}\right)\frac{\partial k}{\partial x_j}\right] + G_k - \rho\varepsilon, \tag{2.6}$$

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \frac{\partial(\rho \overline{u}_j \varepsilon)}{\partial x_j} = \frac{\partial}{\partial x_j}\left[\left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\frac{\partial \varepsilon}{\partial x_j}\right] + C_{1\varepsilon}\frac{\varepsilon}{k}G_k - C_{2\varepsilon}\rho\frac{\varepsilon^2}{k}, \tag{2.7}$$

In these equations, $G_k = \mu_t S^2$ represents the generation of turbulence kinetic energy due to mean velocity gradients, where $S = \sqrt{2S_{ij}S_{ij}}$ is the modulus of the mean strain-rate tensor. The turbulent eddy viscosity is $\mu_t = \rho C_\mu k^2/\varepsilon$. The symbols σ_k and σ_ε are the turbulent Prandtl numbers governing diffusion of k and ε, respectively; $C_{1\varepsilon}$ and $C_{2\varepsilon}$ are model constants. The full ANSYS formulation also includes a buoyancy production term G_b and a compressible dilatation term Y_M; both vanish for the incompressible, isothermal water flow simulated here and are therefore omitted [3]. The standard model constants [15] are $C_\mu = 0.09$, $C_{1\varepsilon} = 1.44$, $C_{2\varepsilon} = 1.92$, $\sigma_k = 1.0$, $\sigma_\varepsilon = 1.3$.

The *realizable* k-ε model [3,25] retains the same k-equation but introduces two modifications to the ε transport equation: C_μ becomes a variable function of mean strain and rotation rates rather than a constant (ensuring physical realizability of normal stresses), and the ε production and destruction terms are reformulated. The realizable model uses different constants: $C_2 = 1.9$ and $\sigma_\varepsilon = 1.2$ [3, 25]. These changes improve predictions in flows with strong adverse pressure gradients, rotation, and streamline curvature, all of which are present on the stern of a semi-planing hull at $Fr = 0.796$. The realizable variant is selected for this work.

2.3.2 The k-ω SST Model

For comparison, the k-ω SST model [17] blends the k-ω formulation (which integrates accurately through the viscous sublayer) near walls with a k-ε formulation in the free stream (which is less sensitive to inlet conditions). It further limits the eddy viscosity in adverse pressure gradient regions. SST is widely regarded as the current best practice for flows with significant separation. However, when used in low-y^+ mode for full boundary-layer integration, it requires $y^+ \approx 1$ at the first cell, which substantially increases the near-wall cell count and computational cost compared to a wall-function approach.

Table 2.1: Summary comparison of turbulence models for marine CFD.

Model	Near-wall mesh	Strengths	Limitations
Standard k-ε	$y^+ > 30$ (WF)	Robust, cheap	Poor in APG, rotation
Realizable k-ε	$y^+ > 30$ (WF)	Better APG, jets, curvature	Boussinesq limitation
k-ω SST	$y^+ \approx 1$ (low-y^+) or WF	Best overall accuracy	Higher cost, finer mesh

2.4 Near-Wall Treatment and Wall Functions

2.4.1 Structure of the Near-Wall Region

The turbulent boundary layer near a smooth wall exhibits a universal structure described in terms of the dimensionless wall distance and velocity:

$$y^+ = \frac{u_\tau\, y}{\nu}, \qquad u^+ = \frac{\overline{u}}{u_\tau}, \qquad u_\tau = \sqrt{\frac{\tau_w}{\rho}}, \tag{2.8}$$

where τ_w is the wall shear stress and u_τ is the friction velocity. Three regions are identified [20]:

- Viscous sublayer ($y^+ < 5$): molecular viscosity dominates, $u^+ = y^+$,
- Buffer layer ($5 < y^+ < 30$): transitional, neither linear nor logarithmic law holds,
- Log layer ($30 < y^+ < 300$): turbulent stress dominates, $u^+ = \frac{1}{\kappa}\ln(E\, y^+)$. Classical Coles constants are $\kappa = 0.41$ and $E = 9.793$ [20]; ANSYS Fluent [3] uses $\kappa = 0.4187$ and $E = 9.793$.

In a wall-resolved approach, the mesh is refined until $y^+ \approx 1$ in the viscous sublayer and the entire profile is computed. This is extremely expensive at the Reynolds numbers encountered in ship hydrodynamics ($Re \sim 10^7$–10^9): the boundary layer thickness δ scales as $\delta/L \sim Re^{-1/5}$ [20], requiring millions of near-wall cells. Wall functions bridge the gap between the wall and the first computed cell, replacing the near-wall region with analytical profiles.

2.4.2 Standard Wall Functions

Standard wall functions [15] assume the first cell centroid lies in the log layer and impose analytical expressions for wall shear stress, turbulent kinetic energy, and dissipation rate derived from log-law velocity profiles. These relations are exact only within the log layer. When $y_P^+ < 11.225$ (the intersection of the linear sublayer law

and the log law), the log-law assumption breaks down, leading to a spuriously low wall shear stress and incorrect skin friction drag prediction.

2.4.3 Scalable Wall Functions

Scalable wall functions, introduced by Grotjans and Menter and implemented in ANSYS Fluent [3,7], address the buffer-layer degradation by substituting the physical y_P^+ with a clipped effective value y_P^* in all wall function expressions:

$$y_P^* = \max\Big(y_P^+,\ y_{\text{lim}}^+\Big), \qquad y_{\text{lim}}^+ = 11.225. \tag{2.9}$$

The value $y_{\text{lim}}^+ = 11.225$ is the point at which $u^+ = y^+$ (viscous sublayer) and $u^+ = \kappa^{-1}\ln(E\,y^+)$ (log law) intersect. By substituting y_P^* for y_P^+ in the wall function expressions, the wall function is effectively prevented from being applied inside the viscous sublayer: when a cell centroid falls below $y^+ = 11.225$, the model behaves as though the centroid is exactly at the sublayer-log layer boundary, ensuring numerically stable and physically bounded behaviour.

The modified expressions become:

$$\tau_w = \rho\, C_\mu^{1/4}\, k_P^{1/2}\, \frac{\overline{u}_P}{\kappa^{-1}\ln(E\,y_P^*)}, \tag{2.10}$$

$$\varepsilon_P = \frac{C_\mu^{3/4}\, k_P^{3/2}}{\kappa\, y_P^*}. \tag{2.11}$$

In practice, scalable wall functions remain reliable for $30 \le y^+ \le 300$ (the target log-layer range), and they degrade gracefully rather than catastrophically when cells fall below $y^+ = 30$. This robustness is the primary reason for their use in the present simulation. In a VOF multiphase setup, the hull surface is partially submerged and partially in contact with air. Since the free-surface position and spray extent cannot be determined a priori, the near-wall mesh must be refined over the entire hull surface to ensure adequate y^+ in the wetted region. This inevitably produces small first-cell heights above the waterline as well, where the local wall shear stress is near zero, driving y^+ well below 11.225 in those cells. Standard wall functions would produce a spurious wall shear stress in those cells; scalable wall functions clip y^* to y_{lim}^+ and yield a bounded, stable result.

2.4.4 Wall Function Limitations in Adverse Pressure Gradients

Both standard and scalable wall functions assume that turbulent kinetic energy production and dissipation are in balance at the first cell centroid. This holds in

attached flows with mild pressure gradients, where the boundary layer evolves slowly. In regions of strong adverse pressure gradient, flow separation, or reattachment, this balance breaks down: the decelerating flow disrupts the local equilibrium, and the wall function expressions derived from it become inaccurate. On the stern of a semi-planing hull at $Fr = 0.796$, where the boundary layer encounters a significant APG as the hull rises above the waterline, local equilibrium is questionable.

The consequences for prediction accuracy are: the separation point may be displaced upstream, the recirculation zone (if present) may be underestimated in extent, and the pressure drag contribution from the stern may be slightly underestimated. Studies on similar hull forms report that realizable k-ε with scalable wall functions typically overpredicts the attached flow region by a small margin, leading to a total drag error of 3–5% [7, 14]. This level of accuracy is sufficient for comparative optimization: all evaluations use identical wall function settings, so the relative drag ranking between hull variants is preserved even if the absolute drag value carries a systematic offset.

2.4.5 Model and Near-Wall Treatment Selection

Three considerations motivate the selection of the realizable k-ε model with scalable wall functions. First, the optimization requires a large number of CFD evaluations; computational efficiency is therefore a primary concern. Second, the hull geometry changes between evaluations: the boundary layer mesh must remain consistent without manual adjustment of near-wall refinement, which scalable wall functions accommodate more robustly than a wall-integrated approach. Third, comparative studies on ship resistance prediction report that realizable k-ε typically predicts total resistance within 3–5% of experimental data, sufficient to discriminate between hull geometries [14]. The selection follows the recommendations of Eca et al. [7] for optimization-oriented campaigns.

For the present hull at $L_{wl} = 7\,\mathrm{m}$, $V = 6.6\,\mathrm{m/s}$, $Re \approx 4.6 \times 10^7$, the flat-plate friction coefficient from the ITTC 1957 line is:

$$C_F = \frac{0.075}{(\log_{10} Re - 2)^2} = \frac{0.075}{(7.663 - 2)^2} \approx 2.34 \times 10^{-3}, \tag{2.12}$$

giving a friction velocity $u_\tau = \overline{u}\sqrt{C_F/2} \approx 0.226\,\mathrm{m/s}$. The first-cell height $\Delta y_1 = 1\,\mathrm{mm}$ then yields:

$$y^+ = \frac{u_\tau\,\Delta y_1}{\nu} = \frac{0.226 \times 0.001}{10^{-6}} \approx 226, \tag{2.13}$$

placing the first cell well within the log layer and within the optimal range for scalable wall functions. The combination of realizable k-ε with scalable wall functions is therefore retained for all simulations in this work.

2.5 Scale Effects and Full-Scale Strategy

Physical model testing at scale ratio $\lambda = L_m/L_s$ faces an inherent conflict: enforcing Froude similarity ($V_m = V_s\sqrt{\lambda}$) violates Reynolds similarity, yielding a model Reynolds number approximately $\lambda^{3/2}$ times lower than full scale. For $\lambda = 1/10$, $Re_m \approx Re_s/32$: the boundary layer is proportionally thicker, friction drag is larger relative to wave drag, and the flow may be partially laminar.

For semi-planing hulls at high Froude number, additional complications arise: the equilibrium trim and wetted area are speed-dependent, so the form factor is not constant and ITTC extrapolation procedures are unreliable; spray formation is affected by Weber-number scaling at model scale; and turbulence stimulators required at low Re_m perturb the pressure distribution. Published guidelines acknowledge these difficulties and recommend full-scale CFD for high-Froude-number applications [27].

The present work bypasses model-scale testing entirely by simulating directly at full scale ($L_{wl} = 7\,\text{m}$, $Re_s \approx 4.6 \times 10^7$). This eliminates extrapolation uncertainties from speed-dependent wetted area, turbulence stimulation, and form factor variation. The full-scale Reynolds number is well within the range where RANS with wall functions is established and validated against experimental data for comparable hull geometries [14]. At $Re_s \approx 4.6 \times 10^7$, the target first-cell height is $\Delta y_1 = 1\,\text{mm}$, keeping the first-cell centroid in the log layer as required by the scalable wall functions (Section 2.4).

2.6 Free-Surface Modelling: Volume of Fluid Method

The simulation of a hull advancing at speed involves two fluid phases: water below the waterline and air above. The Volume of Fluid (VOF) method [3, 9] tracks phase composition through a scalar volume fraction field $\alpha_w \in [0, 1]$, where $\alpha_w = 1$ is fully immersed, $\alpha_w = 0$ fully in air, and $0 < \alpha_w < 1$ contains the interface. The transport equation for α_w is:

$$\frac{\partial \alpha_w}{\partial t} + \nabla \cdot (\alpha_w \, \mathbf{u}) = 0. \tag{2.14}$$

The RANS equations are solved for a single effective fluid whose density and viscosity are:

$$\rho = \alpha_w \, \rho_w + (1 - \alpha_w) \, \rho_a, \qquad \mu = \alpha_w \, \mu_w + (1 - \alpha_w) \, \mu_a. \tag{2.15}$$

The free surface is defined as the isosurface $\alpha_w = 0.5$.

Standard cell-face interpolation of α_w introduces numerical diffusion that smears the water–air interface over many cells. Fluent's compressive scheme [3] counteracts this by reconstructing face values with a bias toward the high-α_w side, keeping the transition zone within approximately 2–3 cells and preserving a sharp interface.

Surface tension is neglected: the Weber number $We = \rho V^2 L/\sigma \approx 4.3 \times 10^6$ (with $\sigma \approx 0.073\,\mathrm{N/m}$ the surface tension coefficient of water) confirms that inertial forces dominate over surface tension forces by several orders of magnitude at this scale.

Chapter 3

Numerical Setup

3.1 Hull Geometry and Parameterization

3.1.1 Vessel Description

The vessel under study is a 7-metre electric yacht designed by Lumen Yachts for semi-planing operation. The key hull parameters are summarized in Table 3.1. The hull has a clearly defined *chine*: a longitudinal edge separating the relatively flat bottom from the more vertical side panels. In a semi-planing hull, the chine plays a critical hydrodynamic role by delimiting the wetted area, promoting hydrodynamic lift, and controlling flow separation at the hull-waterline junction.

Table 3.1: Key vessel parameters used in all simulations.

Symbol	Parameter	Value	Unit
L_{wl}	Waterline length	7.0	m
m	Mass	1000	kg
W	Weight $= mg$	9810	N
∇	Displaced volume	0.976	m^3
ρ	Water density	1025	kg/m^3
V	Operating speed	6.6	m/s
Fr	Froude number	0.796	–
Re	Reynolds number	4.6×10^7	–
q	Dynamic pressure $= \frac{1}{2}\rho V^2$	22 331	Pa

The operating speed of $V = 6.6\,\mathrm{m/s}$ corresponds to twice the theoretical hull speed. Hull speed is defined as the speed at which the waterline length equals the wavelength of a surface wave, i.e. the wavelength $\lambda = 2\pi V^2/g$ equals L_{wl}:

$$V_{\text{hull}} = \sqrt{\frac{g\,L_{wl}}{2\pi}} = \sqrt{\frac{9.81 \times 7}{2\pi}} \approx 3.3\,\mathrm{m/s}. \tag{3.1}$$

Twice hull speed is therefore $2\,V_{\text{hull}} \approx 6.6\,\text{m/s}$, and the corresponding Froude number is:

$$Fr = \frac{2\,V_{\text{hull}}}{\sqrt{g\,L_{wl}}} = \frac{2}{\sqrt{2\pi}} \approx 0.796. \tag{3.2}$$

Figure 3.1 shows the Grasshopper parametric model of the hull.

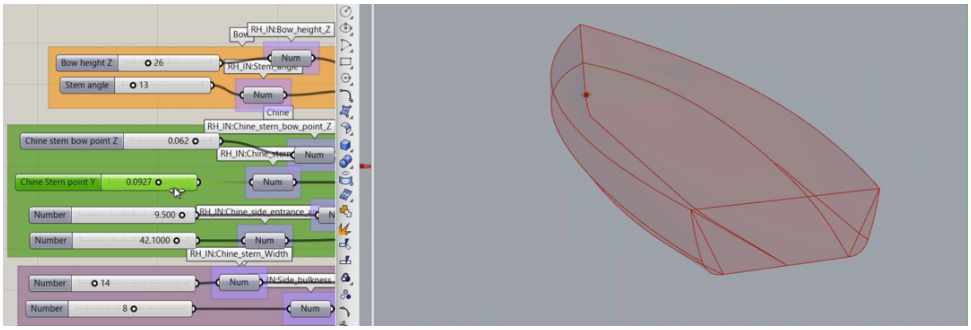


Figure 3.1: Grasshopper parametric model of the Lumen Yachts hull. Left: a portion of the Grasshopper canvas showing a selection of the parameter sliders. Right: the resulting hull geometry displayed in the Rhinoceros viewport. The parametric definition exposes 29 geometric variables controlling the hull shape; `Chine_stern_point_Y` (the lateral chine position at the stern) is the optimization parameter. The full parameter list is given in Appendix F.

3.1.2 Parametric Hull Generation with Grasshopper

The hull geometry is generated through a Grasshopper definition [22], a visual programming environment within Rhinoceros 3D. The definition encodes the hull form as a function of 29 geometric parameters controlling the hull geometry through key control points, curves, and angles. The parametric hull form was provided by Lumen Yachts; the contribution of this work is the closure of the hull surface (transom and deck) required for watertight CFD meshing, and the automation layer connecting the parametric model to the CFD pipeline. Parameters include keel depth and curvature (`Keel_middle_x`, `Keel_bow_X`), stem angle (`Stem_angle`), and mid-body form parameters. The chine is parameterized through three dedicated variables: `Chine_stern_point_Y` (lateral position at the stern, the optimization parameter), `Chine_side_entrance_angle` (entry angle at the bow), and `Chine_side_exit_angle` (exit angle at the stern). The Grasshopper definition generates a complete, closed hull surface including transom, which is a prerequisite for watertight CFD meshing. This closure was implemented specifically for this project and required non-trivial modifications to the original Grasshopper definition.

The hull generation is driven programmatically through Rhino.Compute [23], a server-side API that exposes Rhinoceros geometry operations over HTTP. This allows Python scripts to submit parameter values and receive geometry files without a graphical Rhino session, which is essential for automation.

3.2 Computational Domain

A half-domain is used, exploiting the port-starboard symmetry of the hull and its straight-ahead motion. The fluid domain is constructed around the hull using the FreeCAD Python API and subsequently processed in ANSYS SpaceClaim for zone labelling and format conversion.

Domain extents are specified as multiples of the waterline length L_{wl}:

- $2\,L_{wl}$ upstream of the bow (inlet),
- $4\,L_{wl}$ downstream of the stern (outlet),
- $2\,L_{wl}$ from the centreline to the side (symmetry and far-field),
- $1.5\,L_{wl}$ below the waterline (bottom),
- $0.5\,L_{wl}$ above the waterline (top / air domain).

For $L_{wl} = 7\,\mathrm{m}$, this gives a domain of approximately $49 \times 14 \times 14\ \mathrm{m}^3$ (length $\times$ width $\times$ height). The generous upstream extent ensures that the inlet boundary condition does not interact with the bow wave system.

A body of influence (BOI) is placed around the free-surface region, extending from the inlet to the outlet in the streamwise direction and from the symmetry plane to the far-field wall in the lateral direction. Its extent is sized to encompass the full wave deformation around the hull, serving three purposes: it ensures that wave crests and troughs are captured on an adequately refined mesh, it provides a smooth free-surface interface throughout the domain, and it avoids abrupt mesh transitions near the inlet. Figure 3.2 shows the domain cross-section with the BOI visible around the free surface.

3.2.1 Boundary Conditions

- **Velocity inlet (upstream face):** Flow enters at $V = 6.6\,\mathrm{m/s}$ with open-channel wave boundary conditions enabled. This imposes the correct free-surface elevation at the inlet, avoiding numerical wave generation at the inflow boundary [3].
- **Pressure outlet (downstream face):** Gauge pressure set to zero; backflow is prevented.
- **Top, bottom and side faces:** Treated as stationary no-slip walls. The quantitative impact on hull resistance is confirmed negligible in Section 5.3.2 via a blockage ratio analysis.

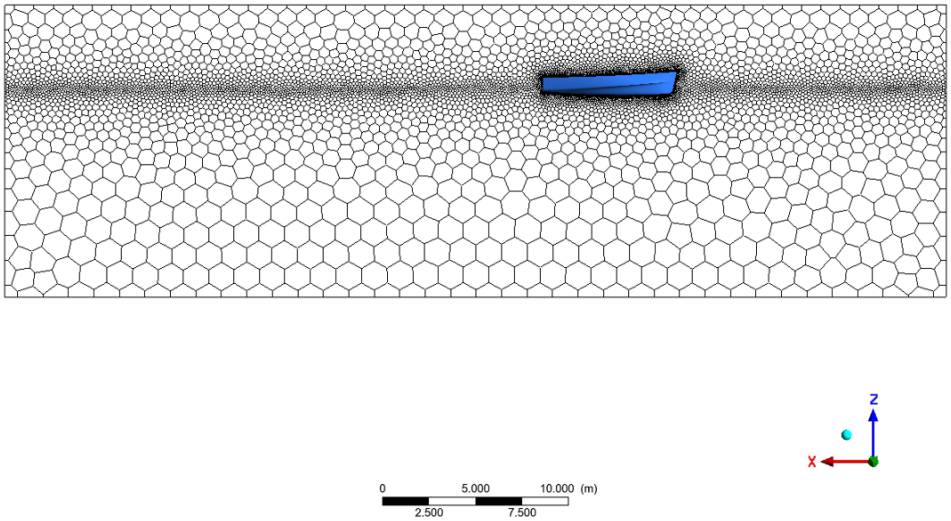


Figure 3.2: Lateral cross-section of the computational domain. The body of influence (BOI) is visible as a refined zone around the free surface. The hull is at the right; the inlet is at the far right, the outlet at the left.

- **Symmetry plane:** Zero normal gradient for all quantities.
- **Hull surface:** No-slip wall with scalable wall function treatment.

The velocity inlet imposes a flat free surface at $z = 0$, with water below and air above, each phase moving at $V = 6.6\,\mathrm{m/s}$ in the positive x-direction. Fluent's Open Channel Wave BC option is used to prevent spurious wave generation at the inflow boundary, which would otherwise introduce numerical waves that propagate into the domain and interact with the hull wave system [3].

Gravity is applied in the negative z-direction with $g = 9.81\,\mathrm{m/s^2}$, and the Volume of Fluid (VOF) method [9] is used with two phases (water and air) to capture the free surface. Open Channel Flow and Open Channel Wave BC options are enabled in Fluent to maintain a realistic air-water interface at the inlet and outlet.

3.3 Mesh Generation

Meshing is performed in ANSYS Fluent 23 using the Watertight Geometry workflow [3]. The workflow is driven entirely by a Python journal file (`.jou`), generated automatically for each hull geometry evaluated during optimization.

The mesh strategy follows these steps:

1. A *body of influence* (BOI) zone, sized at 0.13 m globally, is defined around the hull-domain assembly to provide a locally refined region where free-surface and near-hull flow gradients are highest.

2. A *face size* of 0.03 m (3 cm) is imposed on the hull surface. This corresponds to $\Delta x / L_{wl} = 0.0043$.

3. Sixteen layers of structured boundary layer cells are extruded from the hull surface, with a first-cell height of $\Delta y_1 = 1\,\text{mm}$ and uniform layer growth. This targets $y^+ \approx 226$ (log layer, consistent with scalable wall functions), although local departures from this value occur near the free surface and at the chine, as discussed in Section 5.3.

4. The volume mesh is filled with polyhedral cells, which offer lower numerical diffusion and better orthogonality quality than tetrahedral cells for the same element count.

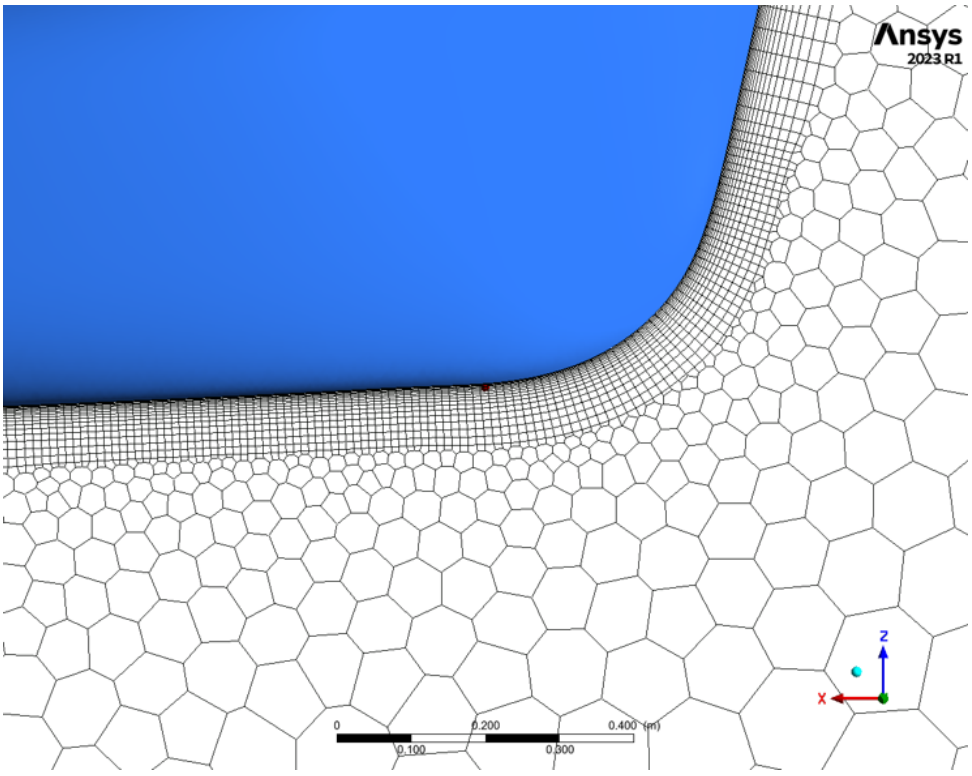


Figure 3.3: Zoomed view of the boundary layer mesh near the hull surface. The 16 structured layers are clearly visible, with the first-cell height of $\Delta y_1 = 1\,\text{mm}$ at the wall, transitioning to polyhedral cells in the far field.

The 3 cm hull face size is retained for all 10 optimization evaluations, yielding approximately 967 000 polyhedral cells per run; the choice is validated by the mesh sensitivity study in Chapter 5.

3.4 Solver Settings

The RANS solver in ANSYS Fluent is run in steady-state mode with the following settings:

- **Turbulence model:** Realizable k-ε with scalable wall functions.
- **Free-surface model:** Volume of Fluid (VOF), two phases (water/air), open-channel flow mode.
- **Phase interaction:** Surface tension neglected (Weber number much greater than unity at this scale).
- **Pressure-velocity coupling:** Coupled scheme [3], which solves the pressure and velocity equations simultaneously rather than sequentially. The coupled scheme offers superior convergence for free-surface flows and complex geometry compared to the segregated SIMPLE algorithm.
- **Spatial discretization:** Second-order upwind for momentum; first-order upwind for turbulent kinetic energy k and dissipation rate ε; PRESTO! for pressure; compressive scheme for the VOF volume fraction.
- **Convergence criteria:** Scaled residuals below 10^{-3} for all equations, plus monitoring of drag, lift, and pitching moment to confirm asymptotic plateau convergence.

The choice of first-order upwind for k and ε is deliberate: these equations are inherently more diffusive than the momentum equation, and second-order discretization tends to produce negative values of k near sharp geometric features, causing solver divergence. Using first-order discretization for the turbulence scalars is standard practice for ship CFD [3, 14].

3.5 Hull Position: Trim and Heave Equilibrium

A freely floating hull at speed assumes a trim angle α and a heave displacement z determined by the balance of hydrodynamic forces and moments with gravity. Figure 3.4 illustrates the forces and moments acting on the hull as modelled in ANSYS Fluent, together with the motor thrust moment that must be accounted for in the equilibrium.

At $Fr = 0.796$, trim and heave are not negligible: the dynamic lift modifies the submerged volume and the pressure distribution significantly. The equilibrium conditions are:

$$F_z(\alpha, z) = W \qquad \text{(vertical force balance)}, \tag{3.3}$$

$$M_y(\alpha, z) = 0 \qquad \text{(pitching moment balance)}, \tag{3.4}$$

where F_z is the total vertical force (lift plus buoyancy) and M_y is the pitching moment about the centre of gravity, both computed from the CFD solution.

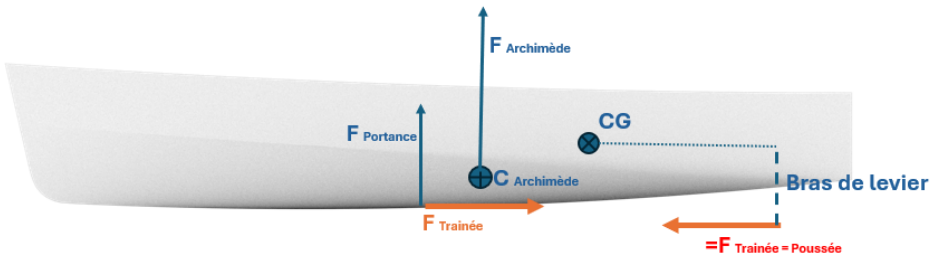


Figure 3.4: Forces and moments acting on the hull at speed. ANSYS Fluent computes the hydrodynamic drag $F_{Trainée}$ (longitudinal resistance), the vertical hydrodynamic forces $F_{Archimede}$ and $F_{Portance}$ (lift plus buoyancy), and the pitching moment M_y about the centre of gravity. The motor thrust $F_{Poussée}$ applied at the propeller shaft introduces an additional pitching moment $M_{\text{motor}} = F_{Poussée} \cdot d_z$, where d_z is the vertical offset between the thrust line and the centre of gravity.

Finding the (α, z) pair that simultaneously satisfies equations (3.3) and (3.4) requires an iterative procedure. This is detailed in Chapter 4, as it is a core component of the automation pipeline.

3.6 Convergence Monitoring for Individual CFD Evaluations

Each steady-state CFD simulation is monitored through two complementary convergence indicators:

Residual convergence. ANSYS Fluent reports scaled residuals for each transported quantity (continuity, x-y-z momentum, k, ε, and volume fraction). The scaled residual for quantity ϕ at iteration n is:

$$R_\phi^{(n)} = \frac{\sum_{\text{cells}} \left| a_P \, \phi_P^{(n)} - \sum_{\text{nb}} a_{\text{nb}} \, \phi_{\text{nb}}^{(n)} - b_P \right|}{\sum_{\text{cells}} \left| a_P \, \phi_P^{(n)} \right|}, \tag{3.5}$$

where a_P, a_{nb}, b_P are the discretization coefficients. Scaled residuals below 10^{-3} indicate that the iterative correction is less than 0.1% of the current field magnitude, which is the adopted stopping criterion.

Monitor point convergence. Residuals alone can reach 10^{-3} while the drag and lift are still drifting, particularly in free-surface simulations where the VOF interface is still adjusting. To detect false convergence, the drag D, lift F_z, and pitching moment M_y are monitored at every iteration.

In practice, all 10 optimization evaluations reached the 10^{-3} residual threshold; the iteration count at which the plateau is reached varies by geometry, typically between 700 and 1000 iterations. A single steady-state CFD run requires approximately 7529 s (about 2 hours) on the available workstation (Intel Core i7, 16 GB RAM, 16 threads), dominated by the ANSYS Fluent solver.

Chapter 4

Automated Optimization Framework

4.1 Parameter Sensitivity Analysis

4.1.1 Motivation

The Grasshopper hull definition exposes 29 geometric parameters (listed in Appendix F). Running a full multi-dimensional optimization over all of them would require thousands of CFD evaluations, far beyond the available computational budget (each optimization evaluation, including trim equilibrium convergence, requires 11 to 45 sequential CFD runs, representing 22 to 90 hours of wall-clock time on the available workstation). A preliminary screening step is therefore performed to identify which parameters most strongly influence drag, so that resources can be concentrated on the most impactful dimensions.

4.1.2 Pearson Correlation Analysis

A set of 50 steady-state CFD simulations is run with random hull geometries sampled from the parameter space, each at a fixed (non-equilibrium) trim and heave. One limitation of this approach is that it does not capture the coupling between hull geometry and equilibrium trim: a parameter that strongly influences the trim angle (and thus indirectly affects drag) may appear less correlated with drag in the fixed-trim screening than it would be in a fully free-floating evaluation. This caveat applies in particular to parameters controlling longitudinal volume distribution. The Pearson correlation coefficient between each parameter p_i and the resulting drag is computed as:

$$r(p_i, D) = \frac{\sum_{j=1}^{N}(p_{i,j} - \bar{p}_i)(D_j - \bar{D})}{\sqrt{\sum_{j=1}^{N}(p_{i,j} - \bar{p}_i)^2}\sqrt{\sum_{j=1}^{N}(D_j - \bar{D})^2}}, \tag{4.1}$$

where $N = 50$ simulations and bars denote sample means.

Table 4.1 reports the ten parameters with the highest absolute correlation. The results are also illustrated in Figure 4.1.

Table 4.1: Top 10 Grasshopper parameters ranked by absolute Pearson correlation with drag (50 preliminary simulations at fixed trim/heave). Physical interpretations are given for the most influential parameters.

Rank	Parameter	$\lvert r \rvert$	Physical role
1	`Chine_stern_point_Y` (p8)	0.55	Lateral position of chine at stern
2	`Chine_side_entrance_angle` (p10)	0.43	Chine entry angle at bow
3	`C2_middle_point_X_direction` (p26)	0.30	Mid-body form control
4	`Radius_X_start` (p16)	0.26	Bottom-to-side transition radius
5	`Keel_middle_x` (p3)	0.23	Longitudinal volume distribution
6	`C2_middle_point_Z_direction` (p28)	0.22	Mid-body vertical control
7	`PB_Pt2_Y` (p21)	0.21	Volume control
8	`Chine_side_exit_angle` (p9)	0.21	Chine exit angle at stern
9	`Keel_bow_X` (p1)	0.21	Bow keel longitudinal position
10	`Stem_angle` (p6)	0.20	Bow stem inclination angle

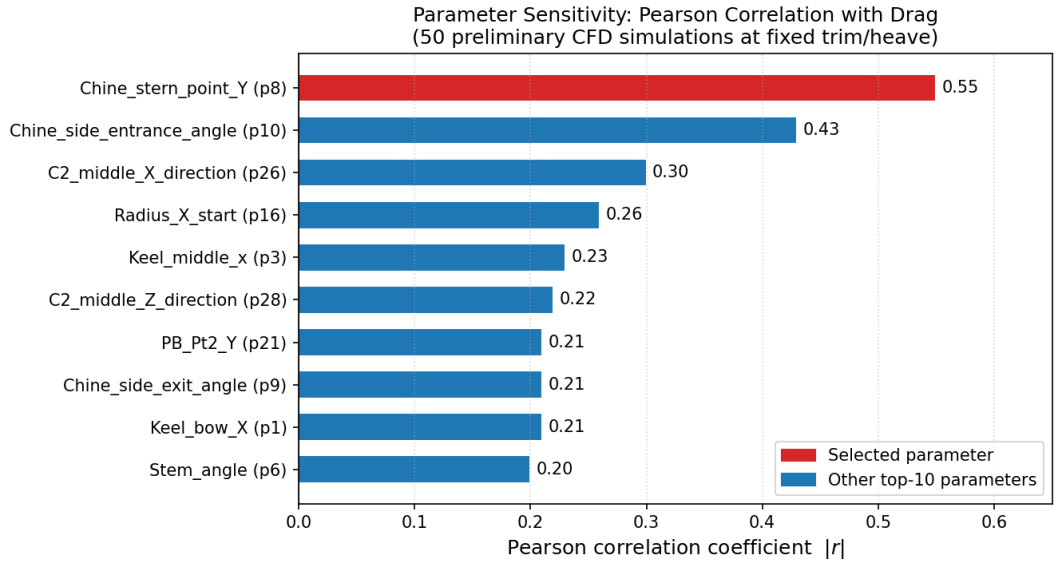


Figure 4.1: Pearson correlation coefficients for the 29 hull parameters. The top 10 are labelled. Chine_stern_point_Y (p8) is clearly dominant.

The stern chine position (Chine_stern_point_Y) emerges as by far the most influential parameter ($r = 0.55$), consistent with its direct effect on the underwater wetted surface at the stern, where hydrodynamic pressure gradients are steepest in the semi-planing regime. Four parameters exceed correlation above 0.25 (up to 0.43), including the optimized parameter, indicating that a multi-parameter optimization targeting this group would ultimately be desirable. The multi-parameter optimizer described in Section 4.4.3 is designed for this purpose.

4.1.3 Choice of Optimization Parameter

For the final optimization campaign, Chine_stern_point_Y is selected as the sole optimization variable. This choice is justified as follows: (i) it has the highest individual correlation with drag; (ii) its physical interpretation is clear, allowing results to be interpreted meaningfully; (iii) limiting the campaign to one parameter allows convergence to be verified within the available computational budget; and (iv) the result demonstrates the end-to-end functionality of the pipeline. A single-parameter optimization provides a clear baseline against which future multi-parameter campaigns can be compared.

4.2 Overview of the Pipeline

The optimization pipeline is fully automated: given a parameter value (or vector of values), the system produces a drag coefficient without any manual intervention. Figure 4.2 illustrates the six stages.

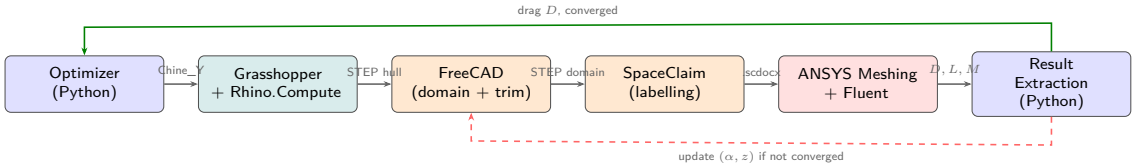


Figure 4.2: Automated optimization pipeline. Each block is one software tool; all transitions are driven by Python scripts. The red dashed arrow shows the trim/heave equilibrium iteration: if force and moment residuals exceed the threshold, the hull position is updated and a new CFD simulation is launched. The green arrow returns the converged drag to the optimizer.

1. **Grasshopper / Rhino.Compute:** The optimization algorithm proposes a value for the chine parameter. A Python script sends this value to Rhino.Compute, which evaluates the Grasshopper definition and returns a closed hull surface in STEP format.

2. **FreeCAD:** A Python FreeCAD script reads the hull STEP file, applies the initial trim angle and heave offset, and constructs the fluid domain solid (Section 3.2). The domain geometry is exported as a STEP file.

3. **ANSYS SpaceClaim:** The domain STEP file is imported, boundary zones are labelled (hull, inlet, outlet, symmetry, etc.), and the watertight geometry is exported in SpaceClaim format (`.scdocx`) for Fluent meshing.

4. **ANSYS Meshing + Fluent:** A Python journal file drives the watertight meshing workflow and the CFD simulation. Drag, lift, pitching moment, and maximum y^+ are written to output files at each iteration.

5. **Trim equilibrium iteration:** The Python trim optimizer reads the CFD output, checks the equilibrium conditions (3.3)–(3.4), and if not satisfied, updates (α, z) and triggers a new FreeCAD/SpaceClaim/Fluent sequence (Section 4.3).

6. **Optimization algorithm:** Once equilibrium is reached, the converged drag value is passed back to the optimizer, which proposes the next parameter value. For the present 1D campaign, the Sobol initialization phase was used throughout (Section 4.4.2); the GP-guided phase was not deployed.

An important advantage of this pipeline is that it functions as a general-purpose optimization tool. Replacing the chine parameter with any other geometric parameter (or set of parameters) requires only a change in the optimizer configuration: the downstream CFD and equilibrium chain is entirely reusable.

4.3 Trim and Heave Equilibrium Algorithm

4.3.1 Equilibrium Solver: Newton-Raphson with Finite-Difference Jacobian

For a given hull shape, trim angle α and heave z are unknowns. The CFD simulation provides $F_z(\alpha, z)$ and $M_y(\alpha, z)$ as outputs, both of which are nonlinear functions of the inputs. The equilibrium system is:

$$\mathbf{R}(\alpha, z) = \begin{pmatrix} F_z(\alpha, z) - W \\ M_y(\alpha, z) \end{pmatrix} = \mathbf{0}. \tag{4.2}$$

The solver uses a Newton-Raphson scheme. At iteration n, the state vector $\mathbf{x}^{(n)} = (\alpha^{(n)}, z^{(n)})$ is updated by:

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \mathbf{J}^{-1}\mathbf{R}(\mathbf{x}^{(n)}), \tag{4.3}$$

where the Jacobian $\mathbf{J} = \partial\mathbf{R}/\partial\mathbf{x}$ is estimated from a linear regression over the last few CFD evaluations. This approach avoids dedicated finite-difference runs: the Jacobian is built from the history of states and residuals accumulated during the iteration.

Several robustness modifications are incorporated in the implementation (version v6):

- **Line search on step size:** the Newton step is tested with factors $[1.0, 0.7, 0.4, 0.2]$ and the first factor that reduces $|\mathbf{R}|$ is accepted.
- **Quasi-convergence memory:** the best state found so far is stored and can be reverted to if the iteration diverges.
- **Warm start:** the trim and heave predicted by a regression model trained on previous hull evaluations provide a starting point for the Newton iteration, substantially reducing the number of CFD calls required per hull.
- **Trim freeze:** when the moment residual is already small, the trim angle is held fixed and only heave is updated, improving convergence stability.

Convergence is declared when $|F_z - W| < 50\,\text{N}$ and $|M_y| < 50\,\text{N}\cdot\text{m}$ simultaneously. In practice, convergence is reached in 11 to 45 CFD evaluations per hull, with an average of 26.2 across the 10 evaluations performed.

4.3.2 Warm Start Regression Model

Each hull evaluation provides, upon completion, a converged (α^*, z^*) pair together with the hull parameter vector $\mathbf{p}$. These data are accumulated across evaluations and used to train a simple linear regression model:

$$\alpha^* \approx \mathbf{a}_\alpha^\top \mathbf{p} + b_\alpha, \tag{4.4}$$

$$z^* \approx \mathbf{a}_z^\top \mathbf{p} + b_z, \tag{4.5}$$

where $\mathbf{a}_\alpha, \mathbf{a}_z$ and b_α, b_z are fit by ordinary least squares on all previous converged evaluations. When a new hull is submitted, the model predicts $(\hat{\alpha}, \hat{z})$ as the warm starting point for the Newton iteration. The benefit is substantial: for evaluation 7 (the best hull, Chine_Y = 0.071), the warm start predicted $\hat{\alpha} = 5.56°$ and $\hat{z} = 165.0\,\text{mm}$, while the equilibrium found was at $\alpha^* = 5.89°$ and $z^* = 156.6\,\text{mm}$, a prediction error of $0.33°$ and $8.4\,\text{mm}$. Starting from a random initial guess instead would have required significantly more Newton iterations to reach convergence.

The warm start model is particularly effective because the equilibrium trim and heave vary smoothly and monotonically with the chine geometry parameter, making them amenable to linear regression even with a small training set.

4.4 Optimization Algorithms

4.4.1 Bayesian Optimization Principles

Bayesian optimization (BO) is a surrogate-based strategy for optimizing expensive black-box functions where gradients are unavailable [24,26]. A probabilistic Gaussian Process (GP) surrogate is built from all evaluations so far and used to select the next evaluation point by maximizing an acquisition function that balances exploration of uncertain regions with exploitation of promising ones. The acquisition function used here is the Lower Confidence Bound (LCB):

$$\mathrm{LCB}(x) = \mu(x) - \beta\,\sigma(x), \tag{4.6}$$

where $\mu(x)$ and $\sigma(x)$ are the GP posterior mean and standard deviation, and β decreases from 2.0 to 0.5 linearly over the optimization. This analytic form is particularly efficient for 1D problems, where the acquisition is maximized over a grid of 500 candidate points.

The GP uses a Matérn 5/2 kernel, which assumes twice-differentiable functions and is physically appropriate for drag as a function of hull geometry.

4.4.2 Initialization: Sobol Sequences

Before the GP is trained, an initialization phase samples the parameter space quasi-uniformly using Sobol sequences [13]. These low-discrepancy sequences provide better coverage of the input space than purely random sampling, ensuring the initial GP is well-conditioned.

Parameter bounds. The exploration range was set to Chine_Y $\in$ [0.055, 0.150]. These bounds are dictated by the Grasshopper parametric definition: below 0.055 or above 0.150, the definition fails to construct a valid spline joining the chine to the hull faces, producing degenerate or non-physical geometries that cannot be meshed.

Adaptive exploration strategy. The campaign was conducted in three successive phases at decreasing tolerance, progressively concentrating resources in the promising region:

- A first, coarse pass at 300 N tolerance spanned the full range [0.055, 0.150] to map the global shape of the drag landscape.
- A second pass at 100 N tolerance refined the picture.
- A final phase at the target 50 N tolerance targeted the region where the minimum was forming.

During the coarser passes it became apparent that drag increases monotonically for Chine_Y > 0.116: no evaluation in the upper portion of the domain yielded a competitive drag value. Consequently, no further evaluations were placed above 0.116; computational resources were reallocated to the lower-Chine_Y region where the minimum was forming. This progressive concentration of evaluations in the promising direction is a practical form of adaptive sampling that complements the quasi-uniform Sobol coverage.

The 10 completed evaluations all belong to this initialization and adaptive refinement phase. The GP-guided exploitation phase was not pursued: the Sobol phase revealed a flat minimum where drag differences are smaller than the numerical noise floor of the trim solver, making further BO evaluations unlikely to identify a meaningfully sharper optimum on this single parameter.

4.4.3 Multi-Parameter Bayesian Optimizer

A multi-parameter Bayesian optimizer has been developed and tested on preliminary runs (up to 5 dimensions). It uses the `BoTorch` library [4] with the q-LogNoisyExpectedImprovement (qLogNEI) acquisition function, which is more suitable than LCB for higher-dimensional problems due to its ability to handle noise and batch proposals. A double convergence criterion is implemented: the optimization terminates when either the maximum pairwise distance between recent candidate proposals falls below a threshold (geometric convergence), or the coefficient of variation of the objective function stabilizes (functional convergence).

This optimizer was not used for the final campaign for two reasons: a 4-parameter Sobol initialization alone would have required $2^4 = 16$ evaluations before any GP-guided phase, exceeding the available budget; and for a single influential parameter, the simpler LCB-based 1D optimizer is both more efficient and more interpretable. The multi-parameter optimizer remains ready for a future campaign targeting the four parameters above correlation above 0.25 from the Pearson analysis.

4.5 Development Path and Abandoned Approaches

The pipeline described above represents the final working implementation. Several alternative approaches were explored and abandoned during development.

4.5.1 GMSH for Volume Mesh Generation

GMSH was initially used as the open-source mesher via its Python API. Three compounding difficulties led to its abandonment. First, the native GMSH mesh format (`.msh`) is not directly supported by Fluent: a custom conversion pipeline

had to be developed, mapping face names from an Abaqus (`.inp`) file that Fluent could export for the same geometry and re-encoding the GMSH output in that format so that boundary conditions could be assigned correctly. Second, generating a structured boundary layer of the required quality proved unreliable with the GMSH scripting interface for these hull geometries. Third, GMSH does not natively produce the polyhedral cell types that the Fluent Watertight Geometry workflow generates. These three issues together made the GMSH path significantly more costly to maintain than switching to Fluent Meshing directly, which resolved all three problems simultaneously and eliminated the format conversion entirely.

4.5.2 Transient Simulation with Free Hull Motion

A transient approach was explored as an alternative to the steady-state trim equilibrium loop. In this formulation, the hull is released as a floating body and allowed to find its own pitch and heave equilibrium under the combined action of hydrodynamic forces and gravity, driven by Fluent's dynamic mesh capability. The approach has two conceptual advantages: it eliminates the outer trim iteration loop entirely, since the equilibrium position emerges naturally from the simulation, and it produces a result free from any force or moment tolerance.

Two practical difficulties led to its abandonment. First, incorporating the motor thrust moment, an offset force applied at the propeller shaft, would have required a user-defined function (UDF) inside Fluent to update the applied moment at each time step, adding significant implementation complexity. Second, and more decisively, the simulation time proved prohibitive. The transient phase during which the hull oscillates before settling to equilibrium is not of engineering interest, yet it must be computed in full. This settling phase can be long, governed by the time scales of hull pitch and heave dynamics and requiring small time steps to satisfy stability (CFL) constraints. Since only the final steady-state resistance is needed, the trim equilibrium algorithm developed here, which targets that state directly, proved considerably faster in practice than waiting for a transient simulation to converge to it.

Chapter 5

Verification and Numerical Assessment

5.1 Mesh Sensitivity Study

5.1.1 Context and Methodology

Grid convergence studies in CFD are traditionally based on Richardson extrapolation [21]: by running systematically refined meshes with a consistent refinement ratio, one estimates the observed order of accuracy p and extrapolates to the zero-mesh-size limit. The Grid Convergence Index (GCI) [6] formalizes this into a numerical uncertainty metric widely recommended by the ITTC [11].

However, this approach has a fundamental limitation when wall functions are used. In standard practice, the first-cell height Δy_1 is kept fixed across mesh levels so that y^+ and the wall treatment remain consistent; only the bulk mesh size Δx is refined. Wall functions replace the numerically resolved near-wall region with an analytical log-law profile, which introduces a near-wall modelling error that remains approximately constant regardless of bulk mesh refinement. The total numerical error is therefore the sum of a bulk discretisation error (which decreases with refinement) and a near-wall modelling error (which does not). The solution does not converge toward the exact PDE solution but toward the wall-function-approximated solution, and the apparent order of convergence p does not reflect the true discretisation accuracy of the bulk scheme. This makes Richardson extrapolation unreliable for wall-function-based simulations [7].

This limitation is not unique to the present work: it affects any wall-function-based RANS simulation. Eca et al. [7] discuss this in detail for ship flow applications and recommend that, for optimization purposes, a qualitative mesh sensitivity assessment (demonstrating that the selected mesh is in the convergent region of the sensitivity curve) is more meaningful than a formal GCI estimate. The present study follows

this recommendation.

5.1.2 Mesh Configurations

Six mesh configurations are evaluated, varying the hull face size from $\Delta x = 1.5\,\text{cm}$ (finest) to $\Delta x = 4\,\text{cm}$ (coarsest), with all other settings (boundary layer, solver, boundary conditions) held constant.

Hydrodynamic forces are presented in dimensionless form following the Froude convention [18], normalizing by the hull weight $W = mg$:

$$C_D = \frac{D}{W}, \qquad C_L = \frac{F_z}{W}, \qquad C_M = \frac{M_y}{W\,L_{wl}}, \tag{5.1}$$

where D, F_z, and M_y are the CFD drag, vertical force, and pitching moment respectively. This normalization is natural for semi-planing vessels, where the dynamic lift F_z approaches W at planing speed: $C_L \to 1$ indicates the hull is fully supported by hydrodynamic lift. The ratio C_D represents the propulsive force required as a fraction of vehicle weight, directly connected to energy consumption and range.

Note that all mesh sensitivity simulations are run at fixed trim and heave corresponding to the base hull at rest ($\alpha = 0°$, $z = 0$): this avoids the confounding effect of geometry-dependent equilibrium trim and allows a fair comparison across mesh levels. The equilibrium trim algorithm is applied only during the optimization campaign (Chapter 4).

Table 5.1 reports the element count, computation time, and percentage deviations from the finest mesh for each configuration. The full table including absolute coefficients is given in Appendix E.

Table 5.1: Mesh sensitivity study: percentage deviation of dimensionless coefficients relative to the finest mesh (1.5 cm reference). The selected mesh ($\Delta x = 3\,\text{cm}$) is highlighted. Base hull at fixed trim ($\alpha = 0°$).

Δx (cm)	Elements	ΔC_D (%)	ΔC_L (%)	ΔC_M (%)	Time (s)
1.5	2 229 372	ref	ref	ref	13807
2.0	1 506 134	+2.11	+2.35	+4.01	8192
2.5	1 329 728	+1.72	+0.09	+0.38	7075
3.0	**967 971**	**+2.71**	+1.35	+2.66	**7529**
3.5	874 897	+5.37	+0.93	+2.06	3889
4.0	812 985	+4.71	+0.91	+2.57	4184

5.1.3 Results and Selected Mesh

The drag coefficient C_D shows a general trend towards the finest mesh (Figure 5.1), but the convergence is not strictly monotonic: the 4 cm mesh gives a slightly lower deviation than the 3.5 cm mesh, and the 2.5 cm mesh lies between the 2 cm and 3 cm meshes. This irregular behaviour is expected: as explained in Section 5.1.1, wall functions impose a fixed near-wall length scale, so the apparent convergence order is low and non-monotonic oscillations are normal. In this context, identifying the "best" mesh from the sensitivity curve alone is not straightforward, and the results should be read as a convergence assessment rather than a precise error estimate.

The 3 cm mesh was selected for the optimization campaign for two practical reasons. First, it produces a relative deviation of 2.71% from the finest computed reference (1.5 cm mesh), which is sufficient to distinguish meaningful drag differences between hull variants while remaining in the convergent part of the sensitivity curve. Second, the 3 cm mesh yields approximately 967 000 elements, within the 1 000 000-element ceiling of the ANSYS student licence, which was used for post-processing and visualization; the next finer mesh (2.5 cm, ≈1 330 000 elements) exceeds this limit and would not be accessible for result inspection in that environment. CFD simulations themselves were performed on the university ANSYS licence, which does not carry this element restriction. The choice of the 3 cm mesh is therefore both technically justified and compatible with the full toolchain.

It should be recalled that these sensitivity results correspond to the base hull at fixed trim ($\alpha = 0°$, $z = 0$). The absolute values of C_D, C_L, and C_M in this study are therefore not directly comparable to the optimization results, which use the full equilibrium trim algorithm.

5.2 Boundary Layer Quality

5.2.1 y^+ Distribution

Figure 5.2 shows the y^+ distribution on the hull surface for the 3 cm mesh at the reference trim and heave. On the wetted hull surface, values reach a maximum of approximately 150, below the flat-plate estimate of $y^+ \approx 226$ derived in Section 2.4. This discrepancy is expected: the ITTC 1957 friction line [18] is calibrated for a flat plate at Reynolds number $Re = V L_{wl}/\nu$, whereas in semi-planing condition the bow rises and the effective wetted length is shorter than L_{wl}, reducing the effective Reynolds number and consequently the local wall shear stress and y^+. In the air domain above the waterline, y^+ drops to approximately 5 (air is less viscous and flows more slowly over the emerged hull). All wetted-surface values lie within the log-layer range appropriate for scalable wall functions ($30 < y^+ < 300$) [3, 7]. No

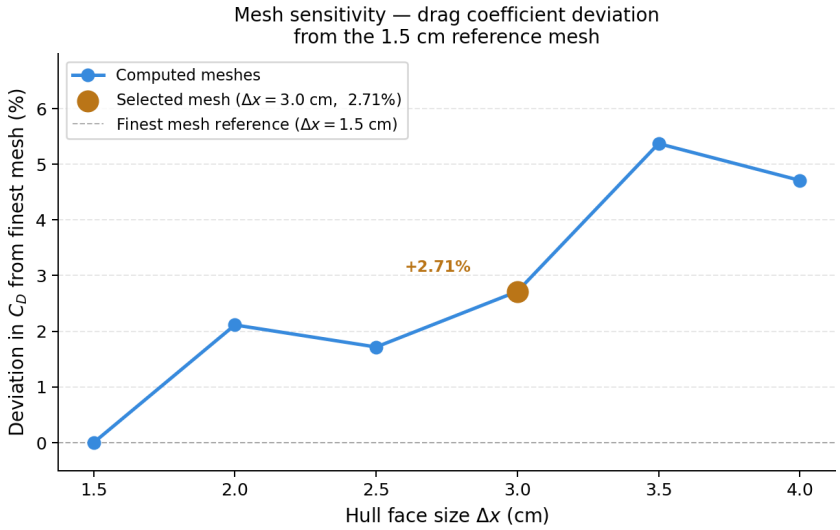


Figure 5.1: Percentage deviation of drag coefficient C_D relative to the finest mesh (1.5 cm reference) as a function of hull face size Δx. The orange marker indicates the selected mesh ($\Delta x = 3$ cm, +2.71% deviation). Richardson extrapolation is not applied: scalable wall functions impose a fixed physical length scale at the wall that does not reduce with bulk mesh refinement, making the observed convergence order unreliable for extrapolation (see Section 5.3).

buffer-layer contamination is observed on the submerged surface.

5.2.2 Boundary Layer Profile: Comparison with the Law of the Wall

The turbulent boundary layer near a smooth wall exhibits a universal structure in inner units. In the viscous sublayer ($y^+ < 5$):

$$u^+ = y^+. \tag{5.2}$$

In the logarithmic layer ($30 \lesssim y^+ \lesssim 300$), the log law holds:

$$u^+ = \frac{1}{\kappa} \ln(y^+) + B, \tag{5.3}$$

with von Kármán constant κ and intercept B. As noted in Section 2.4, ANSYS Fluent uses $\kappa = 0.4187$ and $E = 9.793$ ($B \approx 5.45$) [3]. Classical Coles values are $\kappa = 0.41$, $B \approx 5.56$ [20].

The velocity profile was extracted by tracing lines normal to the hull surface at a mid-ship station, recording the intersection points with the mesh cell centroids to obtain their wall-normal distance Δy and streamwise velocity u; the wall shear stress τ_w was read directly from the hull surface at the same location, giving $u_\tau = \sqrt{\tau_w/\rho}$ and hence $u^+ = u/u_\tau$ and $y^+ = u_\tau \Delta y/\nu$.

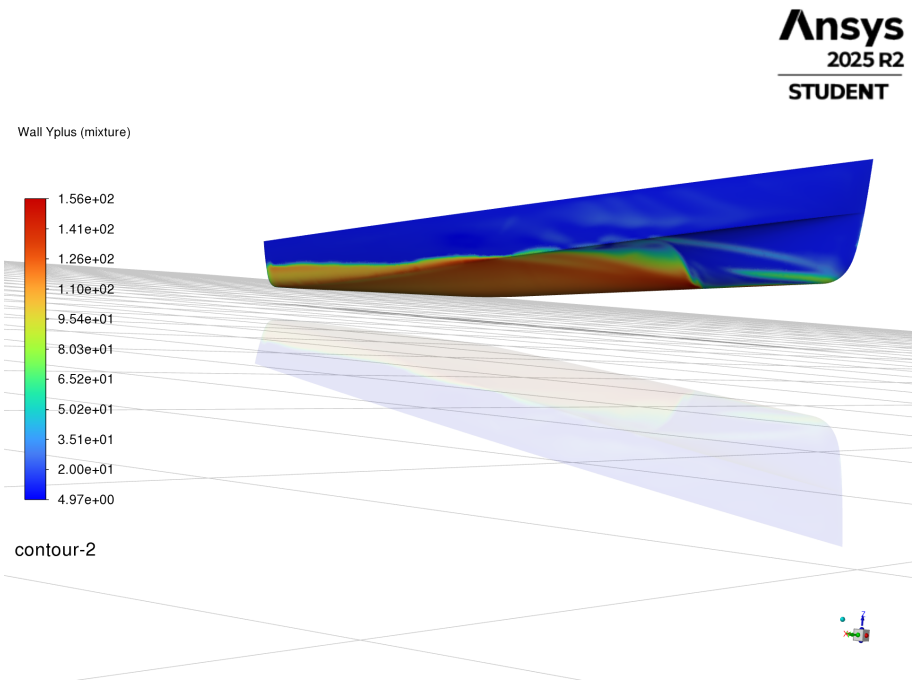


Figure 5.2: y^+ distribution on the hull surface for the selected mesh (3 cm face size). Values on the wetted hull surface reach a maximum of approximately 150; low values ($y^+ \approx 5$) appear in the air domain above the waterline. All submerged-surface values fall within the log-layer range $30 < y^+ < 300$, confirming correct scalable wall function operation.

Figure 5.3 shows the resulting profile. Four data points fall in the log-law region and follow it closely; the three subsequent points deviate above the log law, which is characteristic of the outer (wake) region of the boundary layer where the velocity is increasingly influenced by the free-stream and the log law no longer applies. Interestingly, the best fit to the four log-layer points is obtained with the modern constants $\kappa = 0.383$ and $B = 4.25$, which are considered more accurate than the Coles values for high-Reynolds-number turbulent boundary layers. This is noteworthy since Fluent's wall function uses $\kappa = 0.4187$: the wall function imposes τ_w as a boundary condition, but the velocity field at the cell centroids is determined by the full RANS equations rather than the wall function formula. The resolved RANS velocity profile can therefore follow a log law with constants that differ from those assumed by the wall function, and the modern value $\kappa = 0.383$ appears to describe it more accurately.

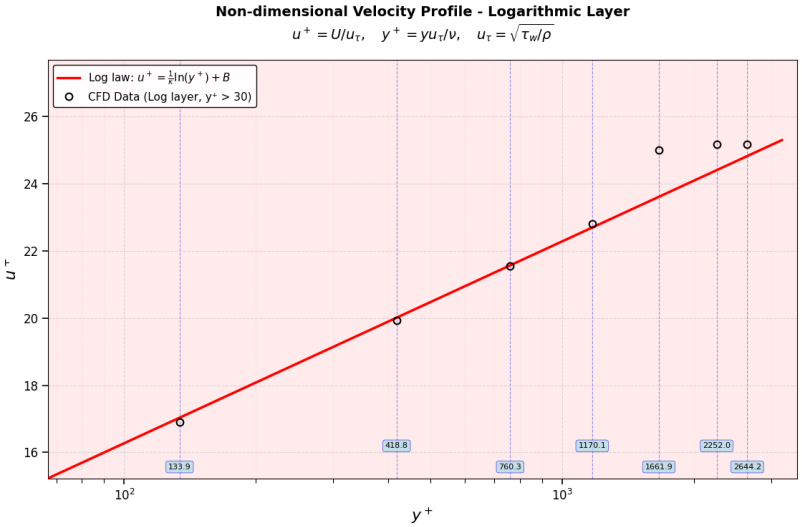


Figure 5.3: Boundary layer velocity profile extracted from CFD (symbols) compared with the log law using modern constants $\kappa = 0.383$, $B = 4.25$. Four points follow the log law; the three outer points depart upward into the wake region.

5.3 Numerical Accuracy and Sensitivity Assessment

5.3.1 Two Distinct Sources of Error: Discretization Imprecision and Model Bias

It is important to distinguish two conceptually different sources of numerical error in this study.

Discretization imprecision (also called numerical precision or consistency) refers to how much the computed drag changes when the mesh is refined, with everything else held constant. The mesh sensitivity results of Table 5.1 quantify this: the 3 cm mesh produces a drag coefficient 2.71% higher than the finest computed mesh (1.5 cm). This is a relative deviation between two points on the sensitivity curve, not an absolute error estimate: the 1.5 cm mesh is itself not the true solution, and the actual discretization error relative to the exact result is unknown. It is a measure of internal consistency between two computations performed under the same physical model, and does not indicate whether either result is close to the true physical value.

This imprecision has a direct consequence for the interpretation of the optimization results. When drag differences between hull variants are smaller than this noise level, the observed ordering between designs cannot be considered physically meaningful: apparent drag variations in that range are at least partly attributable to discretization rather than genuine geometric differences, and the solver cannot reliably rank such designs.

Model bias refers to the systematic deviation of the CFD model from physical reality, regardless of mesh resolution. Even with infinite mesh refinement, the realizable k-ε model with scalable wall functions would still produce a result that may deviate from a towing-tank measurement due to turbulence model assumptions, wall function approximations, and the neglect of spray. This bias is not captured by the mesh sensitivity study at all. No turbulence model sensitivity study was performed; based on published comparisons for similar hull forms [7, 14], replacing realizable k-ε with k-ω SST typically shifts the predicted drag by 2–5%. Since all 10 evaluations use the same model, however, this bias is consistent across hull variants and does not affect the relative drag ranking.

The absolute accuracy of RANS drag predictions for ship-like geometries has been studied through international workshop comparisons. Larsson et al. [14] analyzed the 2010 Gothenburg workshop results: the scatter in drag predictions across different RANS solvers and turbulence models typically spans 3–8% of experimental values. For planing hulls at high Froude numbers, the scatter is generally larger. The 2.71% relative deviation of the 3 cm mesh from the finest computed reference is thus of the same order as the expected physical model bias, which limits the usefulness of a formal GCI estimate. A comparison against towing-tank data for a geometrically similar hull would be required to quantify this model bias; such a validation exercise is identified as a priority for future work (Section 7.3).

From the perspective of optimization, model bias is less critical than the mesh-level deviation: if the bias is approximately constant across hull variants (which is expected, since all hulls share the same topology, speed regime, and turbulence characteristics), then relative drag comparisons remain valid. The expected drag improvement between hull variants in the optimal zone is substantially larger than the 2.71% relative deviation between mesh levels, giving confidence that the drag ranking between geometries with sufficiently different chine positions is reliable.

5.3.2 Sensitivity to Boundary Conditions

The domain extent choice — $2\,L_{wl}$ upstream, $4\,L_{wl}$ downstream, $2\,L_{wl}$ lateral — follows the recommendations of the ITTC CFD guidelines [11] for open-water resistance simulations. The upstream extent is particularly important for semi-planing hulls at high Froude number, where the bow wave system propagates farther upstream than for displacement hulls at moderate speed. No sensitivity study on domain size was performed in the present work, which is a known limitation: domain blockage and wave reflection at the outlet boundary could introduce systematic errors that are not captured by the mesh sensitivity study alone.

Blockage ratio. The impact of the stationary no-slip far-field walls on hull resistance is assessed using the blockage ratio method of Guo et al. [8]. The blockage ratio is defined as:

$$m = \frac{B \times T}{b \times h}, \tag{5.4}$$

where B and T are the full beam and draft of the hull at its inclined equilibrium attitude, and $b \times h$ is the cross-sectional area of the domain. The base hull (evaluation 4, Chine_Y = 0.106) has the largest wetted cross-section and represents the worst case: $B = 2.154\,\mathrm{m}$ (full beam, as measured on the CAD geometry), $T = 0.259\,\mathrm{m}$ (maximum draught of the inclined hull), $b = 28\,\mathrm{m}$ (full-domain equivalent, twice the half-domain width), and $h = 10.5\,\mathrm{m}$ (domain height). This gives:

$$m = \frac{2.154 \times 0.259}{28 \times 10.5} = \frac{0.558}{294} \approx 0.19\%. \tag{5.5}$$

Guo et al. [8] show, for towing-tank configurations with no-slip walls, that blockage effects on resistance are negligible for $m < 0.5\%$ and remain below a 1.5% correction factor even at $m < 1.5\%$. The computed value of 0.19% falls well below the negligibility threshold.

Boundary layer on far-field walls. A complementary argument concerns the boundary layers developing on the no-slip side walls. The relevant length scale is the full domain length $L = 49\,\mathrm{m}$ (the distance over which the wall boundary layer develops), not the hull length. At $V = 6.6\,\mathrm{m/s}$ and kinematic viscosity $\nu = 1.004 \times 10^{-6}\,\mathrm{m^2/s}$, the Reynolds number based on domain length is $Re = 6.6 \times 49/\nu \approx 3.2 \times 10^8$. The Prandtl power-law estimate gives:

$$\delta \approx \frac{0.37\,L}{Re^{0.2}} = \frac{0.37 \times 49}{(3.2 \times 10^8)^{0.2}} \approx 0.36\,\mathrm{m}. \tag{5.6}$$

The lateral clearance between the hull surface and the side wall is $14 - B/2 = 14 - 1.08 \approx 12.9\,\mathrm{m}$, so the far-field boundary layer occupies only $0.36/12.9 \approx 2.8\%$ of this clearance and never interacts with the hull or its wake. Taken together, the low blockage ratio and the thin far-field boundary layers confirm that the no-slip wall treatment has no meaningful influence on the computed hull resistance, and the drag comparisons between hull variants are unaffected by this modelling choice.

The open-channel flow boundary condition at the inlet, which imposes a sharp horizontal free surface, was verified to remain stable over all 1000 solver iterations for all 10 evaluations. No spurious inlet waves were observed in the α_w contour plots, confirming the effectiveness of the open-channel wave BC.

5.4 Comparison with Reference Literature

5.4.1 Order-of-Magnitude Comparison with Molland et al. (2011)

Molland et al. [18] provide resistance data for a range of semi-planing hull series, including the NPL high-speed round-bilge series and Series 62 hard-chine forms, at Froude numbers spanning the semi-planing regime ($Fr \in [0.4, 1.0]$). For hard-chine semi-planing hulls operating around $Fr \approx 0.8$, their data shows total resistance-to-weight ratios D/W typically in the range 0.10–0.18, depending on the beam-to-length ratio, displacement-length ratio, and deadrise angle of the specific hull family.

The present result, $D/W = 1359.8\,\mathrm{N}/9810\,\mathrm{N} = 0.139$ at $Fr = 0.796$, falls within this range, confirming that the absolute level of the CFD predictions is physically plausible for this class and speed regime.

Among the empirical methods compiled in [18], the WUMTIA regression (§10.3.4.2) is the most applicable to the present case: it was developed specifically for hard-chine and round-bilge small craft over $Fr_V \in [0.50, 2.75]$ and $Fr_L \in [0.25, 1.2]$, both satisfied here ($Fr_V = 2.12$, $Fr_L = 0.796$). Applying the C-factor regression for hard-chine hulls [18] with $L/\nabla^{1/3} = 7.06$ and $L/B = 3.25$, interpolated at $Fr_V = 2.12$, yields $C_{\mathrm{FAC}} \approx 64$, an effective power $P_E \approx 6.85\,\mathrm{kW}$, and a total resistance of approximately 1040 N, some 24% below the CFD result.

This gap has a clear methodological origin. First, the WUMTIA database covers vessels of 10–70 m: at $L = 7\,\mathrm{m}$, the present hull falls below the stated validity range and the prediction is an extrapolation. Second, the regression represents an average effective power over a diverse fleet of hull forms, not an optimum. The intrinsic uncertainty of such empirical methods in the semi-planing regime is typically ±15–25% [18]; a 24% gap is therefore within the expected scatter. The comparison confirms that the CFD result is physically consistent with established empirical knowledge while highlighting why numerical simulation, rather than empirical correlation, is required for this specific hull geometry and operating condition [12].

The qualitative physical trends are also consistent with Molland's analysis. He shows that for semi-planing hulls in this Froude number range, resistance is sensitive to chine geometry, deadrise angle, and the degree of transom ventilation, precisely the quantities controlled by Chine_stern_point_Y in the present optimization. The high Pearson correlation ($r = 0.55$) found for this parameter is therefore not surprising from a naval architecture perspective: Molland's data confirm that the chine position at the stern is a primary driver of resistance for this class of hull at $Fr \approx 0.8$.

Chapter 6

Results and Discussion

6.1 Optimization Campaign

6.1.1 Evaluations Performed

Ten hull evaluations were completed within the available computational time. Each evaluation consists of a trim equilibrium iteration (11 to 45 steady-state CFD simulations per hull, 262 total) followed by a converged drag value. CFD drag convergence histories for the best trim iteration of each evaluation are provided in Appendix D; all evaluations reach a stable drag plateau, confirming that the recorded values are not affected by insufficient solver iteration. Figure 6.1 summarises drag, trim, and heave as functions of Chine_stern_point_Y and provides orientation for the discussion that follows; Table 6.1 gives the numerical values.

The parameter range was set to Chine_Y $\in [0.055, 0.150]$ based on geometric feasibility constraints in the Grasshopper definition (values outside this range produce degenerate hull splines). Evaluations were progressively concentrated below 0.116 after early passes showed monotonically increasing drag above that value; the upper portion of the domain received no further resources. The 10 evaluations correspond to the Sobol initialization and adaptive refinement phase, which explores the parameter space quasi-uniformly before concentrating on the promising region. The GP-guided exploitation phase was not pursued: the Sobol phase revealed a flat minimum region where the drag differences between neighbouring hull variants are smaller than the numerical noise floor of the trim solver. A Gaussian Process operating in this regime cannot distinguish between candidate points with statistical confidence, and further evaluations on this single parameter would not identify a meaningfully sharper optimum. The Sobol phase alone is therefore sufficient to characterize the 1D drag landscape.

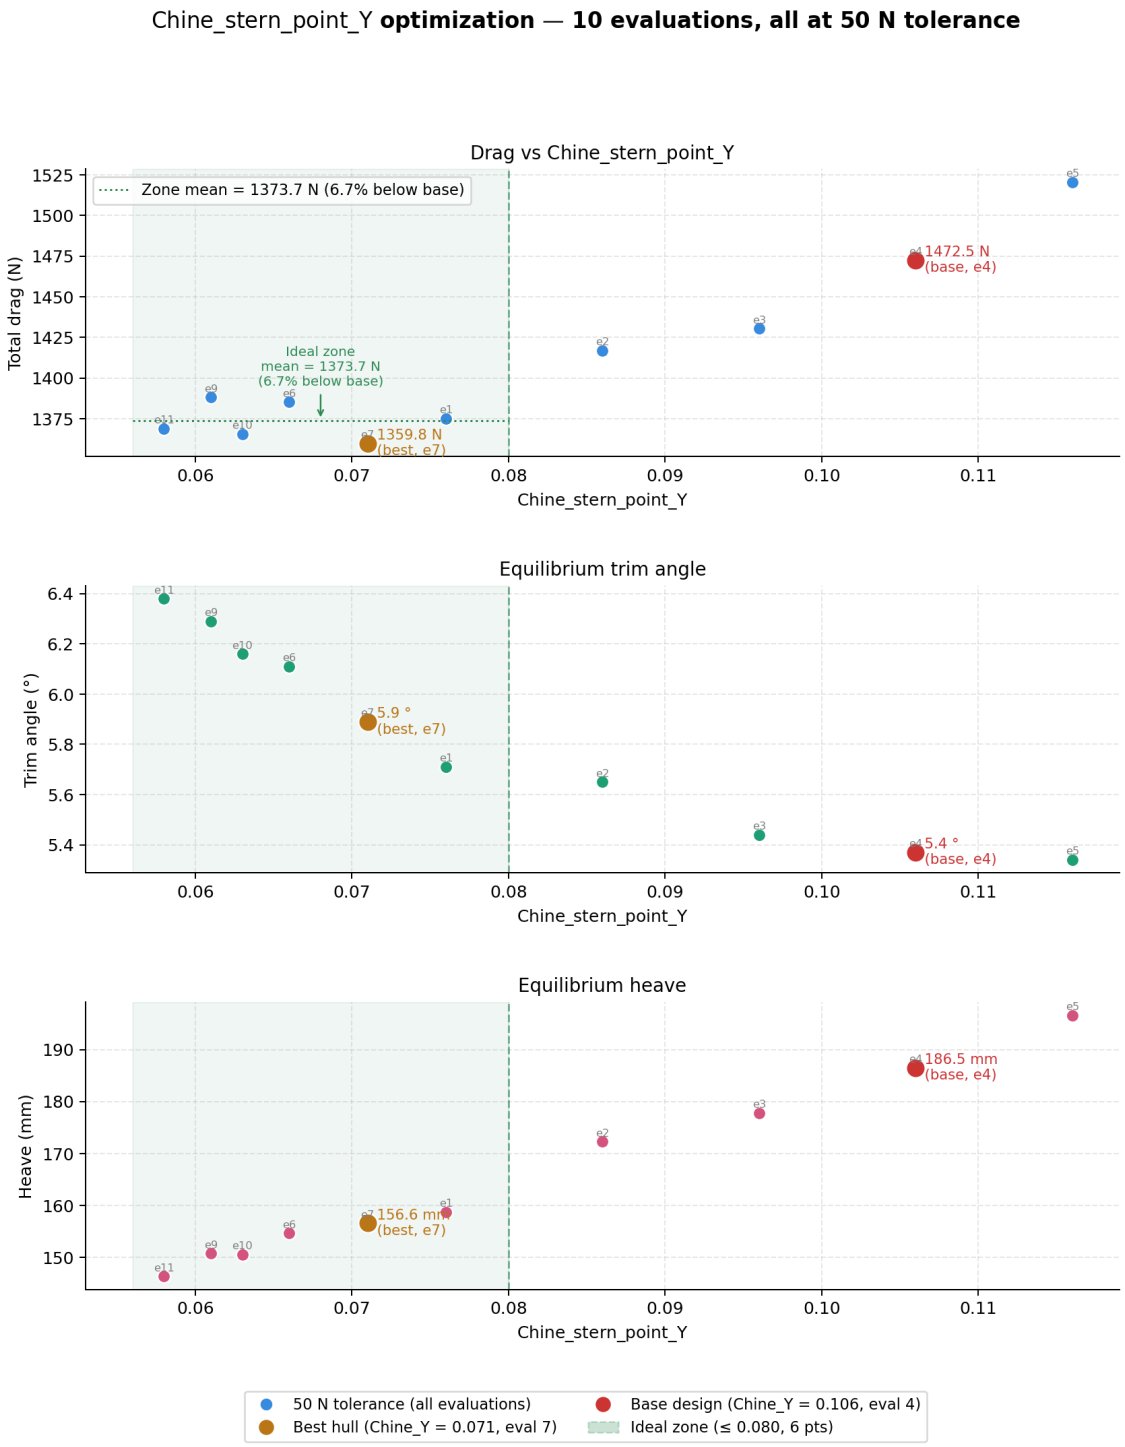


Figure 6.1: Drag force (top), equilibrium trim angle (middle), and equilibrium heave (bottom) as functions of Chine_stern_point_Y. All 10 evaluations converged to the 50 N tolerance (circles). Each point is labelled with its evaluation index (e1–e11 in chronological order; e8 is absent: it was excluded for the reasons detailed in Section 6.2). The best hull (Chine_Y = 0.071, $D = 1359.8$ N, evaluation 7) is shown in orange; the base design (Chine_Y = 0.106, $D = 1472.5$ N, evaluation 4) in red. The green dashed box marks the ideal zone (Chine_Y ≤ 0.080, 6 evaluations), with a mean drag of 1373.7 N, which is 6.7% below the base design. The same zone is highlighted on the trim and heave subplots.

Table 6.1: All 10 optimization evaluations, sorted by Chine_Y (ascending). The best result (lowest drag) is highlighted in bold; the base design (Chine_Y = 0.106) is in italics. All evaluations converged to the 50 N / 50 N·m tolerance. n_{CFD} is the number of steady-state CFD simulations run for that evaluation. The evaluation index is the chronological order of execution, not a rank.

Eval. index	Chine_Y	Drag (N)	α (°)	z (mm)	n_{CFD}
11	0.058	1368.8	6.38	146.4	11
9	0.061	1388.0	6.29	150.8	30
10	0.063	1365.6	6.16	150.6	12
6	0.066	1385.2	6.11	154.7	30
7	**0.071**	**1359.8**	**5.89**	**156.6**	**14**
1	0.076	1375.0	5.71	158.8	17
2	0.086	1417.0	5.65	172.3	26
3	0.096	1430.3	5.44	177.8	38
4	*0.106*	*1472.5*	*5.37*	*186.5*	*45*
5	0.116	1520.5	5.34	196.6	39

6.1.2 Drag as a Function of the Chine Parameter

The drag landscape, shown in Figure 6.1 at the opening of this chapter, exhibits a clearly non-monotonic shape with a pronounced minimum near Chine_Y = 0.071. Three regions are identified:

- **Ideal zone (**Chine_Y ≤ 0.080**):** the six evaluations in this range form a region of low, flat drag. Their mean drag is $\overline{D}_{\mathrm{zone}} = 1373.7\,\mathrm{N}$, which is 6.7% below the base design. Within this zone, the minimum observed is $D = 1359.8\,\mathrm{N}$ at Chine_Y = 0.071, only 1.0% below the zone mean. The spread across all six points is 28.2 N, well within the estimated 2.7% mesh discretization uncertainty: the drag differences between these six geometries are not statistically significant at the present numerical resolution. Any design with Chine_Y $\in [0.058, 0.080]$ can therefore be considered optimal within the limits of this study.

- **Best point found by the solver (**Chine_Y = 0.071**):** lowest drag observed, $D = 1359.8\,\mathrm{N}$. In this region, the chine width at the stern achieves the best balance between generating sufficient hydrodynamic lift and limiting wetted area.

- **Wide chine (**Chine_Y > 0.080**):** drag increases monotonically with Chine_Y. The widest chine evaluated (Chine_Y = 0.116) produces the highest drag observed (1520.5 N, an increase of 11.8% over the optimum). A wider stern chine increases both wetted area and wave-making in this regime.

The trim angle decreases as Chine_Y increases (from approximately 6.4° at Chine_Y = 0.058 to 5.2° at Chine_Y = 0.116), and heave increases correspondingly.

This is physically consistent: a wider chine generates more lift per unit of hull immersion, allowing the hull to ride higher with a smaller trim angle. A regression over the 10 evaluations confirms these visual trends: heave increases linearly with Chine_Y ($R^2 = 0.992$), while the trim angle follows a quadratic trend ($R^2 = 0.987$), flattening toward the edge of the explored range. The contrast is physically plausible: heave is governed by the vertical force balance. A wider chine increases the submerged volume at the stern; to restore equilibrium, the hull rises in proportion, producing a near-linear heave response. Trim, by contrast, is governed by the moment balance, which depends not only on the magnitude of the vertical force but also on where it is applied along the hull. As the chine position changes, the centre of pressure shifts in a way that is not simply proportional to the geometric change, introducing additional curvature into the trim response.

6.2 Trim Equilibrium Summary

6.2.1 Statistics Across All Evaluations

Table 6.2 summarizes the convergence statistics for all 10 evaluations at the final tolerance of 50 N / 50 N · m. The number of CFD calls per evaluation (n_{CFD}) varies from 11 (evaluation 11) to 45 (evaluation 4, the base design, which proved difficult to converge). The total of 262 CFD simulations to obtain 10 equilibrium drag values represents a substantial computational investment; without the warm start, this number would have been significantly higher.

Table 6.2: Trim equilibrium convergence statistics for all 10 evaluations.

Eval. index	α^* (°)	z^* (mm)	n_{CFD}	Warm start	Tol. reached
1	5.71	158.8	17	No	50 N
2	5.65	172.3	26	Yes	50 N
3	5.44	177.8	38	Yes	50 N
4	5.37	186.5	45	Yes	50 N
5	5.34	196.6	39	Yes	50 N
6	6.11	154.7	30	Yes	50 N
7	5.89	156.6	14	Yes	50 N
9	6.29	150.8	30	Yes	50 N
10	6.16	150.6	12	Yes	50 N
11	6.38	146.4	11	Yes	50 N
Mean			26.2		

All ten evaluations converged to the 50 N / 50 N · m criterion. Evaluation 8 (Chine_Y = 0.088) is excluded from all analysis and does not appear in Figure 6.1 or Table 6.1. It was proposed by the Bayesian optimizer and landed close to evaluation 2

(Chine_Y = 0.086), making it largely redundant with an existing data point. When it became necessary to re-run several evaluations to tighten the convergence tolerance from 100 N to 50 N, the author chose not to allocate the corresponding computational resources (approximately two days of wall-clock time) to evaluation 8, given its proximity to evaluation 2 and the low additional information it would have provided. It therefore remains at two trim iterations only and is not representative of a converged drag value.

Figure 6.2 provides a visual summary of the final convergence state for each evaluation: the magnitude of the force and moment residuals at the last trim iteration. All evaluations satisfy both criteria, with the residuals distributed well below the threshold.

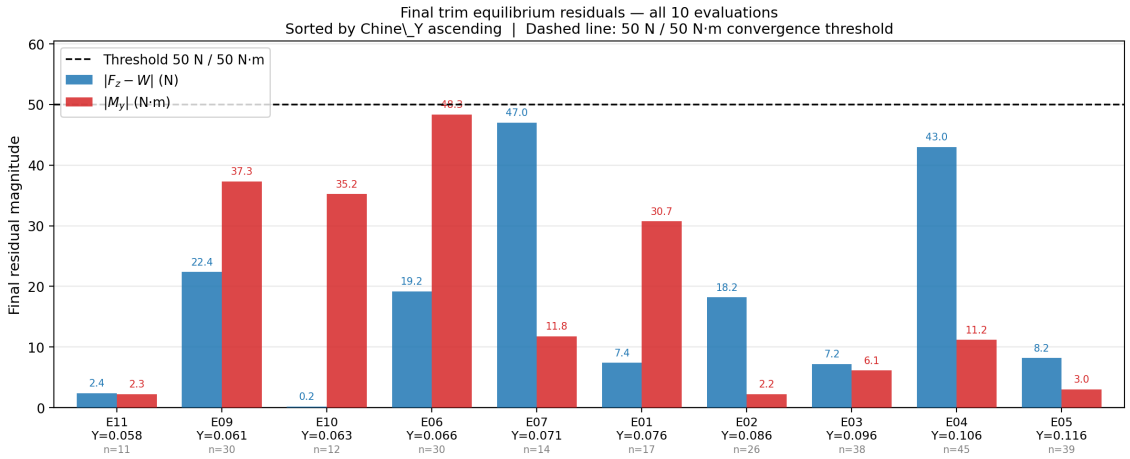


Figure 6.2: Final trim equilibrium residuals for all 10 evaluations, sorted by Chine_Y ascending. Blue bars: force residual $|F_z - W|$ (N). Red bars: pitching moment residual $|M_y|$ (N·m). The dashed line marks the 50 N / 50 N·m convergence threshold. The number of trim iterations n is indicated below each bar group. All evaluations satisfy both criteria simultaneously.

Tightening the tolerance below 50 N was found to be counter-productive. The drag force computed by Fluent fluctuates by approximately ±20–50 N between successive solver iterations due to pressure oscillations near the free surface. Consequently, the Newton-Raphson Jacobian, estimated from the residual change between consecutive CFD evaluations, becomes unreliable when the target tolerance falls below this numerical noise floor. Attempting to push below 50 N causes the optimizer to chase numerical noise, consuming many extra CFD evaluations without any meaningful improvement in trim accuracy.

6.2.2 Example Convergence: Evaluation 7

Figure 6.3 shows the convergence of the trim equilibrium for evaluation 7 (Chine_Y = 0.071, the best hull found). The warm start predicted $\alpha = 5.56°$ and $z = 165.0$ mm.

The Newton-Raphson iteration corrects the trim angle progressively, navigating through states with large force residuals (above 900 N at iteration 1) as the hull is initially far from equilibrium. After each CFD evaluation, the Jacobian is updated and the step size is adjusted by line search. By iteration 14 both residuals are below 50 N, satisfying the convergence criterion. The converged drag is $D = 1359.8$ N.

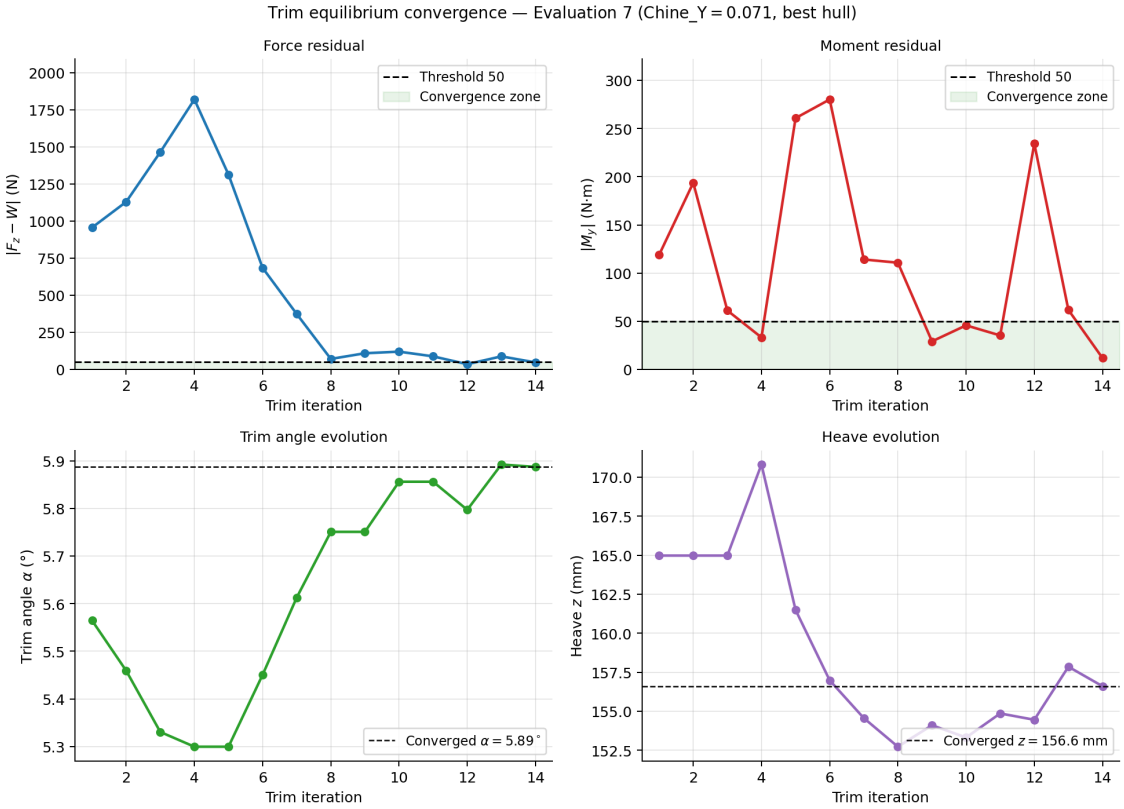


Figure 6.3: Trim equilibrium convergence for evaluation 7 (Chine_Y = 0.071, best hull). Top row: force residual $|F_z - W|$ (left) and pitching moment residual $|M_y|$ (right); the green band marks the convergence zone below the 50 N / 50 N·m threshold (dashed line). Bottom row: evolution of the trim angle α (left) and heave z (right) at each Newton-Raphson iteration, converging to $\alpha = 5.89°$ and $z = 156.6$ mm by iteration 14.

6.3 Flow Field Analysis at the Optimal Configuration

At $Fr = 0.796$, the VOF simulation captures a bow wave crest, a stern depression, and a turbulent wake, all typical of the semi-planing regime. The velocity magnitude peaks near the bow shoulder ($\approx 1.1\,V$) and decreases towards the stern where the flow separates slightly at the chine; the y^+ distribution (Figure 5.2) follows the same pattern and confirms that scalable wall functions operate within their valid range throughout. A wider stern chine increases the wetted planing area at the stern,

generating more hydrodynamic lift aft of the centre of gravity: this is the primary mechanism behind the decreasing trim angle observed with increasing Chine_Y. The running trim angle results from moment equilibrium about the centre of gravity; a narrower chine produces less lift aft, so the hull must adopt a larger bow-up angle to rebalance the pitching moments, which simultaneously keeps the transom closer to the waterline and reduces transom ventilation. This coupling between chine geometry, trim, and drag underlines the importance of the free-floating equilibrium algorithm: a fixed-trim comparison would systematically misrank hull variants.

Figures 6.4 and 6.5 show the VOF free-surface rendering of the optimal hull at the computed equilibrium trim. These visualisations illustrate the physical realism of the simulation: the free surface deforms continuously around the hull, the bow wave propagates outward and rearward, and the water detaches cleanly from the transom at the stern, leaving the characteristic low-pressure depression behind the hull.

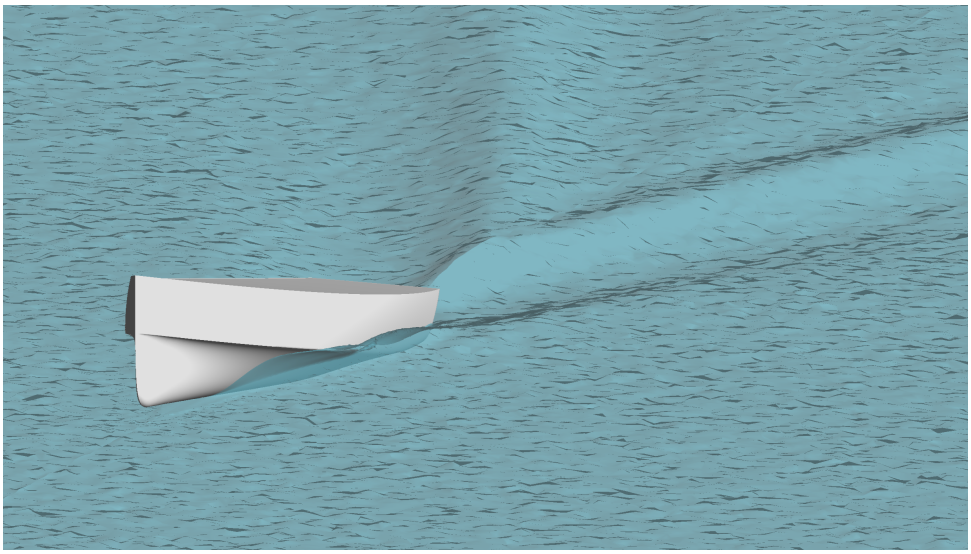


Figure 6.4: VOF free-surface rendering of the optimal hull (eval 7, Chine_Y = 0.071) at $Fr = 0.796$, quarter view from starboard. The bow wave, diverging wake, and clean water detachment at the transom are clearly visible.

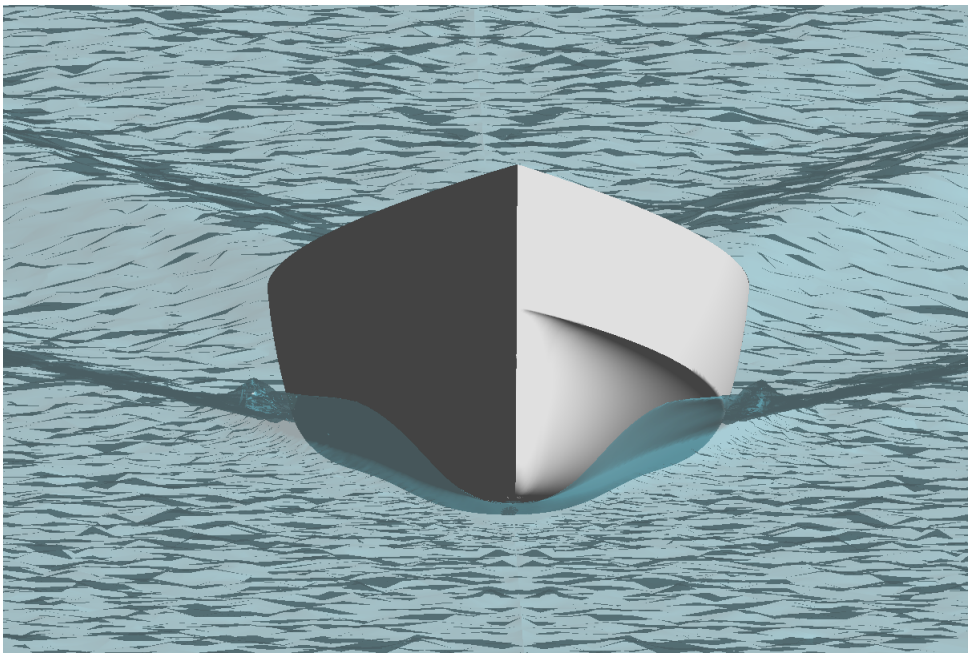


Figure 6.5: Front view of the optimal hull at $Fr = 0.796$. The bow wave pushes water outward on both sides.

6.4 Best Hull Geometry and Comparison with the Base Design

Three flow field quantities illustrate the physical differences between the optimal and base configurations: the free-surface elevation, the wall shear stress distribution, and the static pressure on the hull surface.

6.4.1 Free-Surface Elevation

Figures 6.6 and 6.7 show the free-surface Z topology for both hulls viewed from above. The colour scale represents the water surface elevation in metres relative to the undisturbed free surface.

Both configurations exhibit the bow wave crest and diverging wake pattern characteristic of the semi-planing regime. The base hull shows a marginally higher bow wave crest ($+0.391\,\mathrm{m}$ vs. $+0.383\,\mathrm{m}$), while the optimal hull presents a deeper stern trough ($-0.273\,\mathrm{m}$ vs. $-0.225\,\mathrm{m}$). Part of this difference is explained by the different equilibrium trim: the optimal hull runs at $\alpha_{\mathrm{opt}} = 5.89°$ versus $\alpha_{\mathrm{base}} = 5.37°$, so the absolute vertical position of the hull surface differs and the minimum values are not directly comparable without correcting for the hull reference position. The marginally higher bow wave crest of the base hull is qualitatively consistent with a larger wave-making energy input, which contributes to its higher total drag. The deeper stern trough of the optimal hull is attributed to its larger equilibrium trim

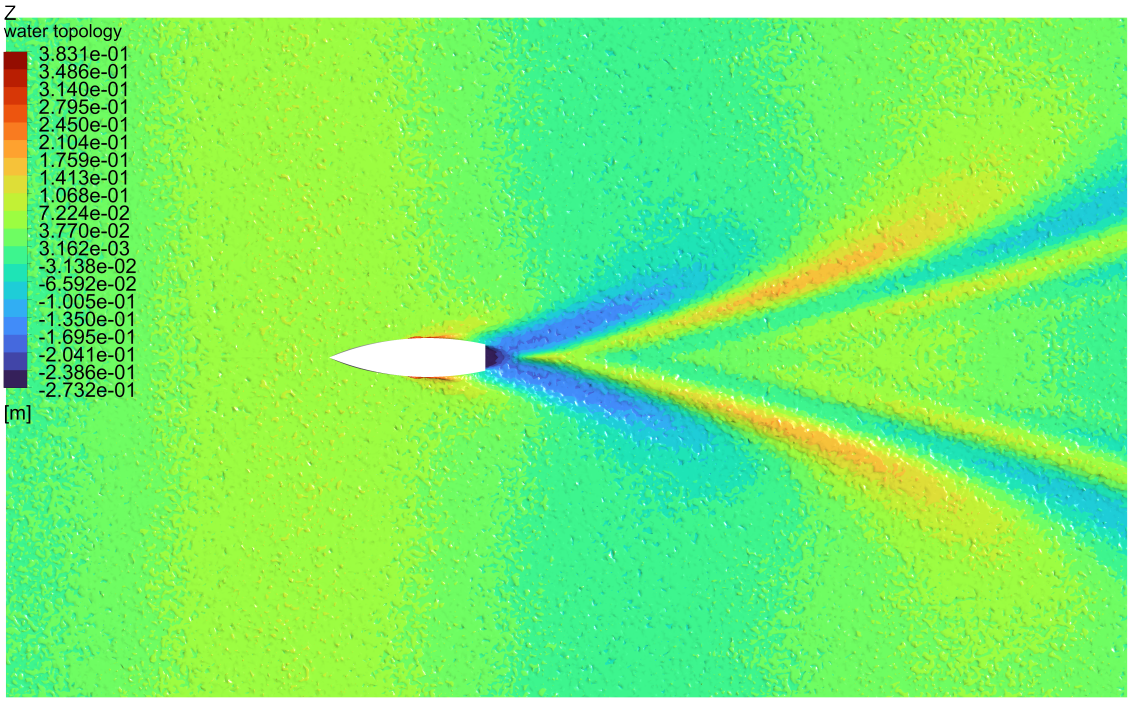


Figure 6.6: Free-surface Z elevation at $Fr = 0.796$ — optimal hull (eval 7, Chine_Y = 0.071): range -0.273 to $+0.383$ m.

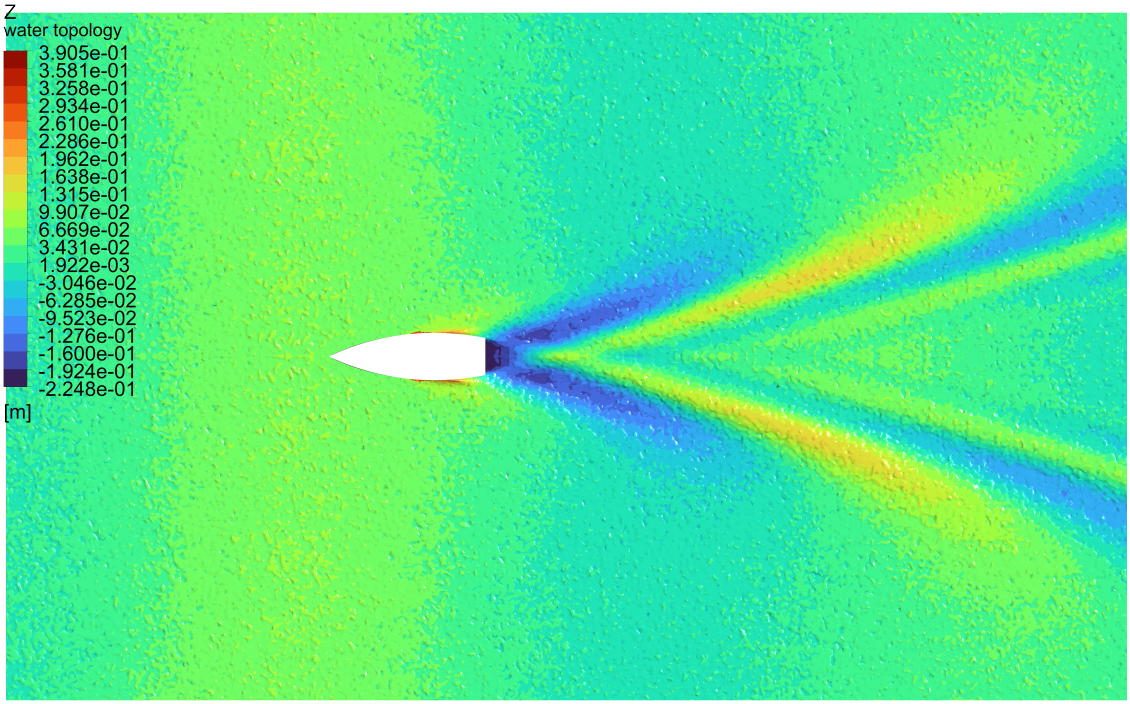
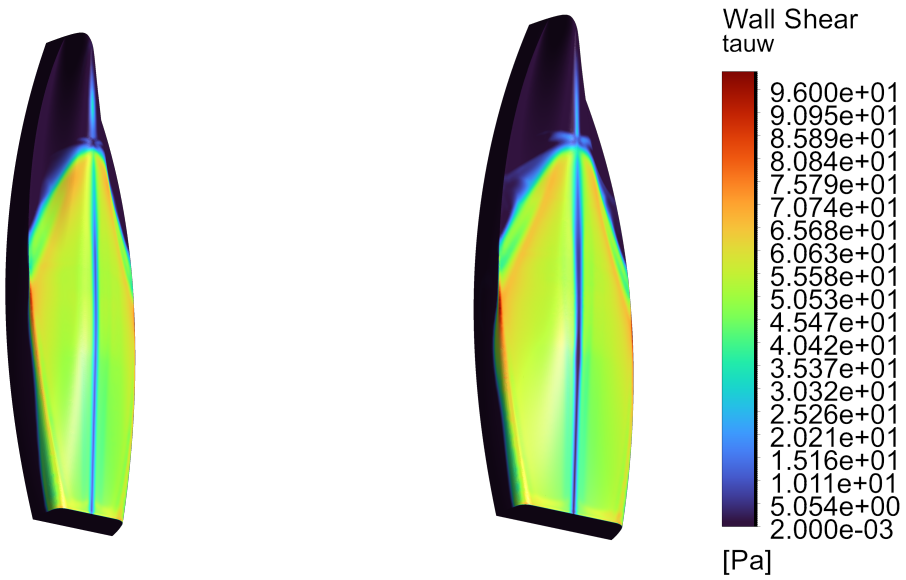


Figure 6.7: Free-surface Z elevation at $Fr = 0.796$ — base hull (eval 4, Chine_Y = 0.106): range -0.225 to $+0.391$ m.

angle (5.89° vs. 5.37°), which lowers the stern relative to the undisturbed waterline, and is therefore not indicative of greater wave-making.

6.4.2 Wall Shear Stress

Figure 6.8 shows the wall shear stress τ_w on the hull surface for both configurations.



(a) Optimal hull (eval 7): peak $\tau_w \approx 88.8\,\text{Pa}$.

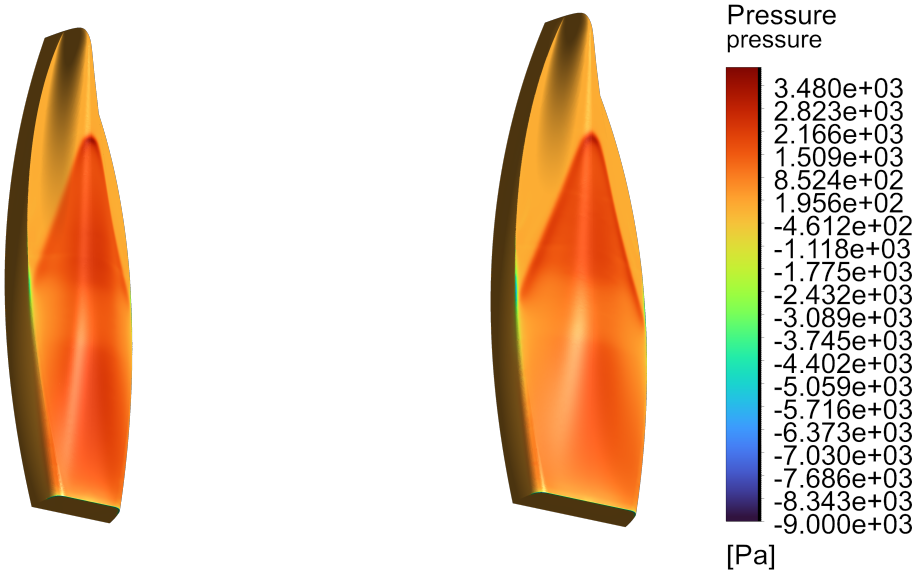
(b) Base hull (eval 4): peak $\tau_w \approx 98.2\,\text{Pa}$.

Figure 6.8: Wall shear stress τ_w [Pa] on the hull surface. Left: optimal configuration. Right: base configuration. The high-shear region is concentrated at the bow shoulder in both cases.

The base hull exhibits a higher peak wall shear stress ($\approx 98.2\,\text{Pa}$ vs. $88.8\,\text{Pa}$ for the optimal). Beyond this localised peak at the bow shoulder, the wider chine of the base hull increases the total wetted area of the planing bottom, raising the integrated viscous force over the entire hull surface. The force decomposition confirms this: the viscous drag of the base hull ($\approx 32.9\%$ of $1472.5\,\text{N} = 484.5\,\text{N}$) exceeds that of the optimal hull ($\approx 30.3\%$ of $1359.8\,\text{N} = 412.0\,\text{N}$) by $72.5\,\text{N}$, which accounts for 64% of the total drag reduction of $112.7\,\text{N}$. Viscous drag is therefore the dominant mechanism behind the improvement.

6.4.3 Pressure Distribution on the Hull Surface

Figure 6.9 compares the static pressure field on the wetted hull surface. Positive values indicate stagnation-like pressure regions; negative values indicate suction.



(a) Optimal hull (eval 7): range −7 427 to +3 558 Pa.

(b) Base hull (eval 4): range −8 682 to +3 099 Pa.

Figure 6.9: Static pressure [Pa] on the hull surface. Left: optimal configuration. Right: base configuration.

The base hull reaches a more negative minimum pressure (−8 682 Pa vs. −7 427 Pa for the optimal), indicating stronger suction regions on its surface that increase form drag. The pressure drag of the base hull (≈ 67.1% of 1472.5 N = 988.0 N) exceeds that of the optimal hull (≈ 69.7% of 1359.8 N = 947.8 N) by 40.2 N, accounting for the remaining 36% of the total drag reduction. Taken together, the two mechanisms fully explain the 112.7 N improvement: viscous drag contributes 64% and pressure drag 36%, with the narrower stern chine acting on both simultaneously by reducing the wetted planing area and moderating the suction field at the stern.

6.4.4 Comparison with the Original Lumen Yachts Hull

Evaluation 4 corresponds to the original stern chine width of the Lumen Yachts hull. The present campaign operates at $Fr = 0.796$ (2× hull speed), a semi-planing regime where the optimal chine geometry differs substantially from that suitable for displacement speed. At the tighter 50 N tolerance, the base design drag is $D_{\text{base}} = 1472.5\,\text{N}$. The minimum-drag hull achieves $D_{\text{opt}} = 1359.8\,\text{N}$, representing a drag reduction of:

$$\Delta D = \frac{D_{\text{base}} - D_{\text{opt}}}{D_{\text{base}}} = \frac{1472.5 - 1359.8}{1472.5} \approx \mathbf{7.7}\%. \tag{6.1}$$

This improvement is achieved by narrowing the stern chine from Chine_Y = 0.106 to Chine_Y = 0.071, a 33% reduction in the chine lateral position. The equilibrium

trim also changes: from $\alpha_{\text{base}} = 5.37°$ (base) to $\alpha_{\text{opt}} = 5.89°$ (optimum), reflecting the different lift-drag trade-off at the narrower chine.

All six evaluations in the ideal zone (Chine_Y ≤ 0.080, Figure 6.1) span only 28.2 N around a mean of 1373.7 N, well within the estimated 2.7% mesh discretization uncertainty. This indicates a *flat minimum*: the individual drag differences within the zone are smaller than the numerical resolution of the simulation, so no single point can be identified as uniquely optimal. All six designs in this range offer a meaningful improvement over the base hull, from -5.7% (eval 9, Chine_Y $= 0.061$) to -7.7% (eval 7, Chine_Y $= 0.071$), with a zone mean of -6.7%.

This improvement in drag has a direct impact on electric range. For a battery of fixed energy, range scales inversely with drag at constant speed ($R \propto D^{-1}$), so the range gain is not equal to the drag reduction but to:

$$\frac{\Delta R}{R_{\text{base}}} = \frac{D_{\text{base}}}{D_{\text{opt}}} - 1 = \frac{1472.5}{1359.8} - 1 \approx 8.3\%. \tag{6.2}$$

The optimized hull therefore extends the cruising range at design speed by approximately 8.3%, a meaningful engineering gain achieved with zero change in hull length or mass.

6.5 Discussion

6.5.1 Physical Interpretation

The flat minimum in drag around Chine_Y ≈ 0.058–0.071 reflects the classical trade-off in semi-planing hull design between wave-making resistance and friction drag. At $Fr = 0.796$, the hull is in the transition between displacement and fully planing regimes. A narrow chine at the stern (small Chine_Y) reduces wetted area and friction, but the stern becomes unable to generate sufficient dynamic lift, the hull sinks deeper (larger displacement), and wave-making resistance increases. A wide chine (large Chine_Y) provides lift but increases friction and generates a stronger stern vortex system, both of which add to total resistance.

Fluent provides a native decomposition of total drag into pressure drag and viscous drag. Post-processing of the two reference configurations (obtained with ANSYS student edition 2025 and expressed as percentages) gives:

- **Optimal hull (eval 7,** Chine_Y $= 0.071$**):** pressure drag $\approx 69.7\%$, viscous drag $\approx 30.3\%$.
- **Base hull (eval 4,** Chine_Y $= 0.106$**):** pressure drag $\approx 67.1\%$, viscous drag $\approx 32.9\%$.

The optimal hull shows a slightly higher pressure-drag fraction and a lower viscous-drag fraction than the base hull. This is consistent with the physical interpretation: the narrower stern chine primarily acts on friction through a reduced wetted area, while the pressure drag also decreases but to a lesser extent. As a result, the pressure-drag fraction increases in relative terms, even though both components decrease in absolute terms. The total drag is nevertheless lower in absolute terms, confirming that the friction reduction dominates.

The strong dependence on trim angle observed in Table 6.1 is also notable. The equilibrium trim varies from 5.36° to 6.38° across the parameter range, a significant spread. This highlights the importance of accounting for free-floating trim in CFD-based optimization: a simulation at a fixed trim would assign incorrect drag values to hull variants that naturally trim differently, potentially misidentifying the optimal geometry.

6.5.2 Comparison to Bretscher (2017)

Bretscher [5] reported drag reductions of 10–15% over his baseline hull through multi-parameter non-linear optimization of a round-bilge displacement craft. The most meaningful comparison with the present campaign is against the Lumen Yachts base hull (Chine_Y = 0.106, eval 4, $D = 1472.5\,\mathrm{N}$), which was itself derived from the Bretscher reference geometry optimized for displacement speed. The minimum-drag configuration (Chine_Y = 0.071, eval 7, $D = 1359.8\,\mathrm{N}$) achieves a **7.7% drag reduction** relative to this baseline — a meaningful result considering that only a single geometric parameter was varied. The improvement relative to the first randomly sampled Sobol point (eval 1, Chine_Y = 0.076, $D = 1375.0\,\mathrm{N}$) is a more modest 1.1%, because this initial point happened to land close to the optimal region by chance. In contrast to Bretscher, the present improvement is achieved within a fully automated pipeline requiring no manual intervention between evaluations, and the framework is immediately deployable for a multi-parameter campaign.

6.5.3 Limitations and Sources of Uncertainty

The key limitations of this work are discussed in Section 5.3 (numerical accuracy) and Chapter 7 (scope and modelling choices). The campaign covers a single geometric parameter out of four with correlation above 0.25; all simulations assume calm water; and the turbulence model (realizable k-ε) may underpredict stern separation compared to SST. The 7.7% drag improvement over the base design is substantially larger than the 2.7% discretization uncertainty, so the optimum region is reliable despite these limitations.

Chapter 7

Conclusion

7.1 Summary and Scientific Perspective

Titled *Hull Optimization for Fast and Efficient Semi-Planing Craft using RANS CFD combined with the Volume of Fluid Method*, this thesis builds and validates a fully automated CFD-based optimization pipeline for a semi-planing hull and uses it to identify the drag-optimal stern chine geometry at $Fr = 0.796$. The pipeline integrates parametric geometry generation (Grasshopper, Rhino.Compute), fluid domain construction (FreeCAD, SpaceClaim), mesh generation and flow simulation (ANSYS Meshing and Fluent), and result extraction and optimization (Python), running from parameter specification to converged drag output without manual intervention.

The automation is the primary scientific contribution. By closing the loop between the parametric geometry tool and the CFD solver through Python orchestration, the per-evaluation overhead (geometry modelling, mesh generation, solver configuration, and result extraction) is fully automated, removing the bottleneck that would otherwise make iterative optimization impractical. The trim equilibrium algorithm, a Newton-Raphson scheme with warm start, line search, and quasi-convergence memory, proved particularly essential: at $Fr = 0.796$, hull variants with different chine geometries trim at different angles. A fixed-trim comparison would systematically misrank them, as confirmed by the spread of equilibrium angles from 5.37° to 6.38° across the parameter range explored. This finding confirms the importance of free-floating equilibrium identified in the literature [16] and extends it to an automated multi-evaluation context.

The Pearson sensitivity study, conducted on 50 preliminary simulations, identified the stern chine position (Chine_stern_point_Y) as the dominant geometric parameter controlling drag at this Froude number, with $r = 0.55$. The subsequent 1D Sobol exploration revealed a flat optimal zone (Chine_Y $\in$ [0.058, 0.080]) with

a mean drag 6.7% below the base design. The flatness of this minimum is itself a physically informative result: within the optimal zone, drag is insensitive to small perturbations in the chine position, implying robustness to manufacturing variation.

A result that deserves explicit discussion is the trade-off between hydrodynamic efficiency and transverse stability. A narrower stern chine reduces drag at cruising speed but also reduces the hydrodynamic righting force. The hull that minimizes resistance is not necessarily the hull that provides the most comfortable or safe sea-keeping behaviour. This tension between fineness and stability is a classical constraint in semi-planing hull design: the chine range identified here (Chine_Y $\in [0.058, 0.080]$) provides the best drag performance, but a final design choice should be validated against stability criteria appropriate for the intended operating conditions.

7.2 Limitations

The primary limitation is the single-parameter scope of the optimization campaign. The Pearson study identified four parameters with correlation above 0.25; optimizing only one leaves most of the design space unexplored. With 262 CFD simulations already performed (including trim equilibrium iterations), the computational investment was substantial for a single parameter; extending to all four parameters requires a proportionally larger effort.

A secondary limitation is the turbulence model: realizable k-ε with scalable wall functions is computationally efficient but less accurate than SST for stern separation. In the absence of experimental towing-tank validation data for this hull, the model bias remains unquantified. Results should be interpreted as comparative numerical trends within the investigated setup, not as fully validated predictions of real-world performance.

Finally, all simulations assume calm water and steady state. The dynamic response of a semi-planing hull in waves involves coupled heave, pitch, and roll motions that are beyond the scope of steady RANS.

7.3 Perspectives

The automated framework is the principal lasting contribution of this thesis. With a larger computational budget, it enables:

- A multi-parameter optimization campaign over the four parameters with correlation above 0.25 identified by the Pearson analysis; the infrastructure, comprising the parametric model, automation pipeline, and multi-parameter Bayesian optimizer, is already in place and tested.

- Extension to free-surface seakeeping simulations using overset mesh techniques, for which the Fluent infrastructure is available.
- Higher-fidelity verification using the SST turbulence model, and quantitative validation against towing-tank measurements for a geometrically comparable hull.
- Transfer of the pipeline to other hull forms under development by Lumen Yachts, since the automation layer is geometry-agnostic and the Grasshopper definition is parametric.

What began as a proof-of-concept for a fast and efficient semi-planing craft has produced something more durable: a shift from manual, sequential simulation to an automated loop where engineering judgment operates at the level of design decisions rather than execution steps. Drag was reduced, the geometry was optimized, and the method was validated on a real engineering problem. The pipeline is ready for the next campaign.

Bibliography

[1] ANSYS Inc. ANSYS Fluent theory guide, release 24.2 — boussinesq approach vs. reynolds stress transport models. https://ansyshelp.ansys.com/Views/Secured/corp/v242/en/flu_th/flu_th_sec_bouss_vs_rst.html, 2024. Accessible with valid ANSYS licence.

[2] ANSYS Inc. ANSYS Fluent theory guide, release 24.2 — reynolds averaging. https://ansyshelp.ansys.com/Views/Secured/corp/v242/en/flu_th/flu_th_sec_reynolds_averaging.html, 2024. Accessible with valid ANSYS licence.

[3] ANSYS Inc. ANSYS Fluent user's guide, release 25.2, 2025.

[4] M. Balandat et al. BoTorch: A framework for efficient monte-carlo Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2020.

[5] E. Bretscher. *Hull Design for Fast and Efficient Round-Bilge Pure-Displacement Craft using RANS CFD and Non-Linear Optimisation*. PhD thesis, University of Auckland, New Zealand, 2017.

[6] I. B. Celik, U. Ghia, P. J. Roache, C. J. Freitas, H. W. Coleman, and P. E. Raad. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering*, 130(7):078001, 2008.

[7] L. Eça, M. Kerkvliet, and S. L. Toxopeus. Estimation of discretization errors in viscous flow simulations based on the RANS equations: Wall-resolved versus wall-functions. *Journal of Verification, Validation and Uncertainty Quantification*, 9:041003, 2024.

[8] Chun-yu Guo, Pei Xu, Chao Wang, and Zi Kan. Numerical and experimental study of blockage effect correction method in towing tank. *China Ocean Engineering*, 33(5):522–536, 2019.

[9] C. W. Hirt and B. D. Nichols. Volume of fluid (vof) method for the dynamics of free boundaries. *Journal of Computational Physics*, 39(1):201–225, 1981.

[10] International Towing Tank Conference. ITTC recommended procedures: Testing and extrapolation methods — high speed marine vehicles resistance test. ITTC Procedure 7.5-02-05-01, 2002.

[11] International Towing Tank Conference. ITTC recommended procedures: Uncertainty analysis in CFD verification and validation. ITTC Procedure 7.5-03-01-01, 2017.

[12] H. Islam, M. Ventura, C. Guedes Soares, M. Tadros, and H. S. Abdelwahab. Comparison between empirical and CFD based methods for ship resistance and power prediction. In C. Guedes Soares and T. A. Santos, editors, *Trends in Maritime Technology and Engineering*. CRC Press, 2022.

[13] S. Joe and F. Y. Kuo. Constructing Sobol' sequences with better two-dimensional projections. *SIAM Journal on Scientific Computing*, 30(5):2635–2654, 2008.

[14] L. Larsson, F. Stern, and M. Visonneau, editors. *Numerical Ship Hydrodynamics: An Assessment of the Gothenburg 2010 Workshop*. Springer, 2014.

[15] B. E. Launder and D. B. Spalding. The numerical computation of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 3(2):269–289, 1974.

[16] A. Le Strat and M. Terziev. Numerical modelling of the effects of trim on the resistance of sailing planing hulls. In *Proceedings of the Postgraduate Maritime Technology Conference (PostGrad MarTec)*, 2024.

[17] F. R. Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8):1598–1605, 1994.

[18] A. F. Molland, S. R. Turnock, and D. A. Hudson. *Ship Resistance and Propulsion*. Cambridge University Press, 2011.

[19] D. Peri and E. F. Campana. Multidisciplinary design optimization of a naval surface combatant. *Journal of Ship Research*, 47(1):1–12, 2003.

[20] S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.

[21] L. F. Richardson. The approximate arithmetical solution by finite differences of physical problems involving differential equations. *Philosophical Transactions of the Royal Society A*, 210:307–357, 1911.

[22] Robert McNeel & Associates. Rhinoceros 3d and grasshopper. https://www.rhino3d.com, 2024.

[23] Robert McNeel & Associates. Rhino.compute. https://developer.rhino3d.com/guides/compute/, 2024.

[24] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[25] T.-H. Shih, W. W. Liou, A. Shabbir, Z. Yang, and J. Zhu. A new k-ε eddy viscosity model for high reynolds number turbulent flows. *Computers & Fluids*, 24(3):227–238, 1995.

[26] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, pages 2960–2968, 2012.

[27] M. Terziev, T. Tezdogan, and A. Incecik. Scale effects and full-scale ship hydrodynamics: A review. *Ocean Engineering*, 245:110496, 2022.

[28] H. K. Versteeg and W. Malalasekera. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*. Pearson Education, 2nd edition, 2007.

Appendix A

Nomenclature and Constants

Table A.1: Symbols and physical constants used throughout this thesis.

Symbol	Definition	Value	Unit
L_{wl}	Waterline length	7.0	m
m	Hull mass	1000	kg
$W = mg$	Hull weight	9810	N
g	Gravitational acceleration	9.81	m/s^2
ρ	Seawater density	1025	kg/m^3
ν	Seawater kinematic viscosity	1.0×10^{-6}	m^2/s
V	Operating speed	6.6	m/s
Fr	Froude number $= V/\sqrt{gL_{wl}}$	0.796	–
Re	Reynolds number $= VL_{wl}/\nu$	4.6×10^7	–
q	Dynamic pressure $= \frac{1}{2}\rho V^2$	22 331	Pa
∇	Displaced volume $= m/\rho$	0.976	m^3
u_τ	Friction velocity $= \sqrt{\tau_w/\rho}$	≈ 0.226	m/s
y^+	Dimensionless wall distance	5–150 (hull surface)	–
κ	Von Kármán constant	0.4187 (Fluent) / 0.41 (Coles)	–
E	Log-law constant	9.793	–
C_μ	Turbulence model constant (realizable)	variable	–
C_μ	Turbulence model constant (standard)	0.09	–
α	Trim angle at equilibrium	5.2–6.4	°
z	Heave at equilibrium	146–197	mm
C_F	Flat-plate friction coefficient (ITTC 1957)	–	–

Appendix B

Optimization Results Summary

Table B.1: Complete optimization results in chronological order (evaluation index). All values correspond to the 50 N convergence criterion. Evaluation 8 is excluded (terminated prematurely, not representative of equilibrium).

Eval	Chine_Y	Drag (N)	α (°)	z (mm)	n_{CFD}
1	0.076	1375.0	5.71	158.8	17
2	0.086	1417.0	5.65	172.3	26
3	0.096	1430.3	5.44	177.8	38
4	0.106	1472.5	5.37	186.5	45
5	0.116	1520.5	5.34	196.6	39
6	0.066	1385.2	6.11	154.7	30
7	0.071	1359.8	5.89	156.6	14
9	0.061	1388.0	6.29	150.8	30
10	0.063	1365.6	6.16	150.6	12
11	0.058	1368.8	6.38	146.4	11
Total CFD simulations					262

Appendix C

Trim Equilibrium Convergence — All Evaluations

The following figures show the trim equilibrium convergence (force residual $|F_z - W|$ and pitching moment residual $|M_y|$) for all 10 optimization evaluations, ordered by Chine_Y ascending. The 50 N / 50 N·m convergence threshold is shown as a dashed line; the green band marks the convergence zone. All evaluations reach both criteria within their allotted trim iterations. Generated by `codes_et_resultats/plot_residus_all_eval`

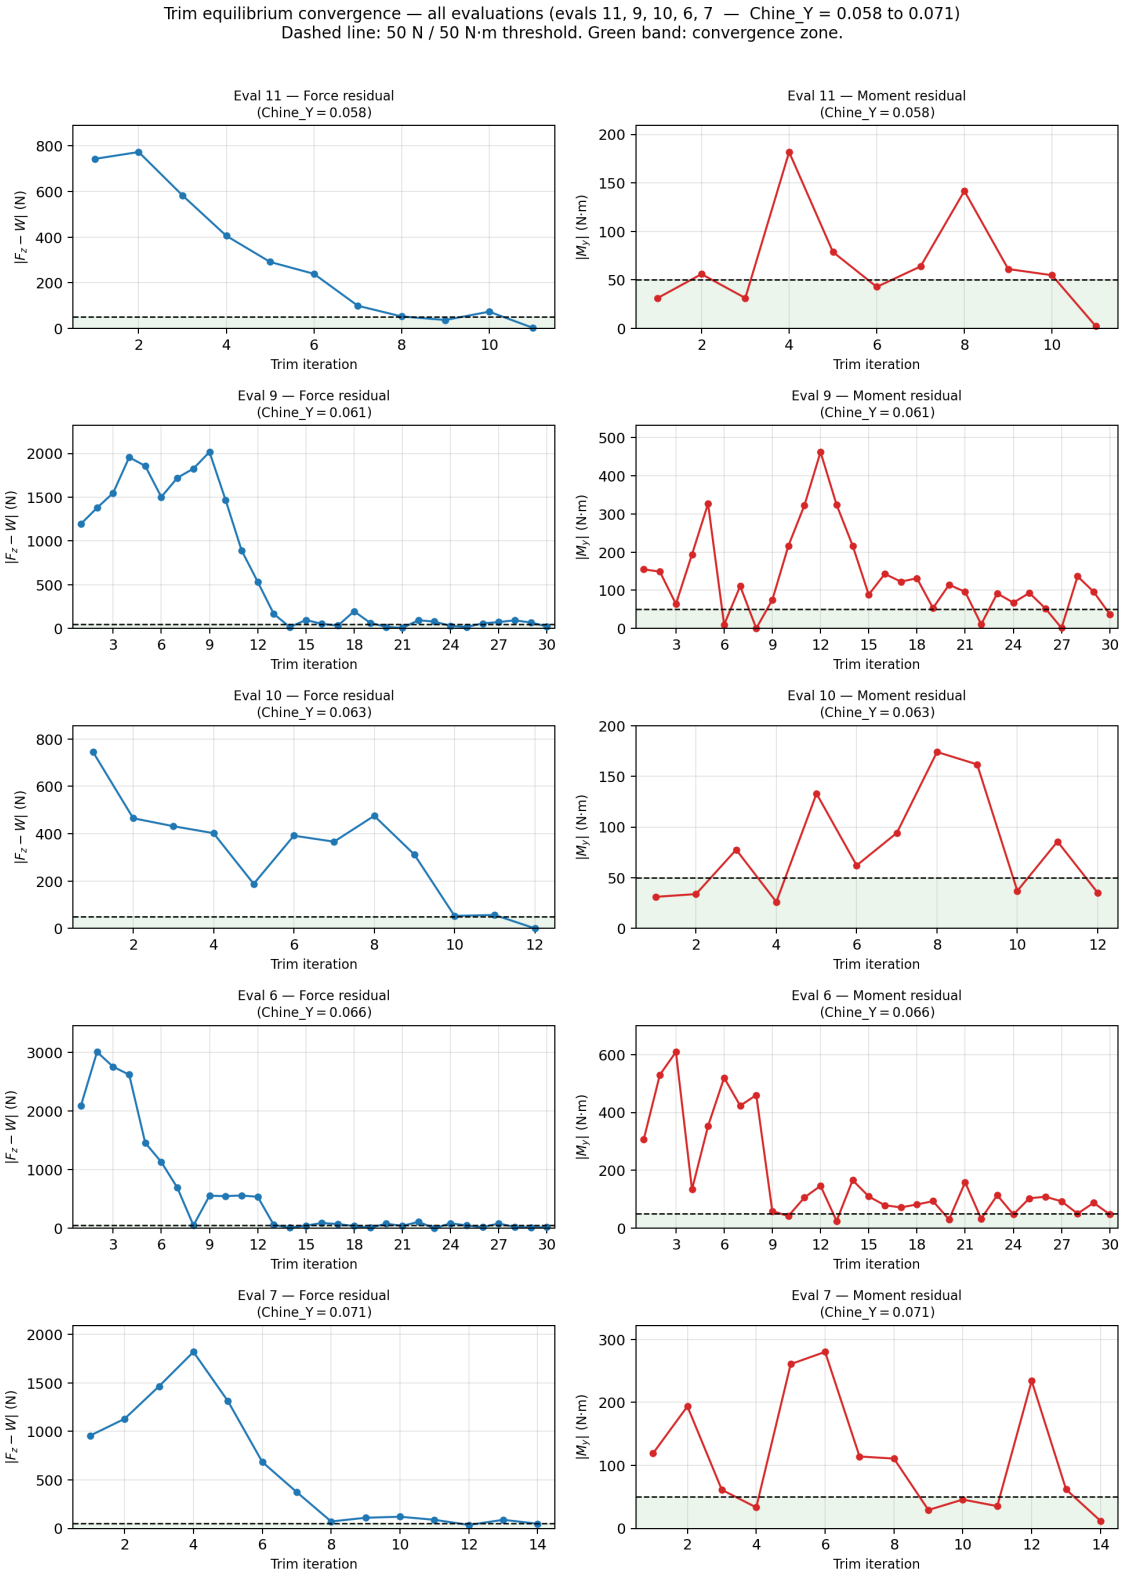


Figure C.1: Trim equilibrium convergence — evaluations 11, 9, 10, 6, 7 (Chine_Y = 0.058 to 0.071). Left column: force residual $|F_z - W|$ (N). Right column: pitching moment residual $|M_y|$ (N·m). Dashed line: 50 N / 50 N·m threshold. Green band: convergence zone.

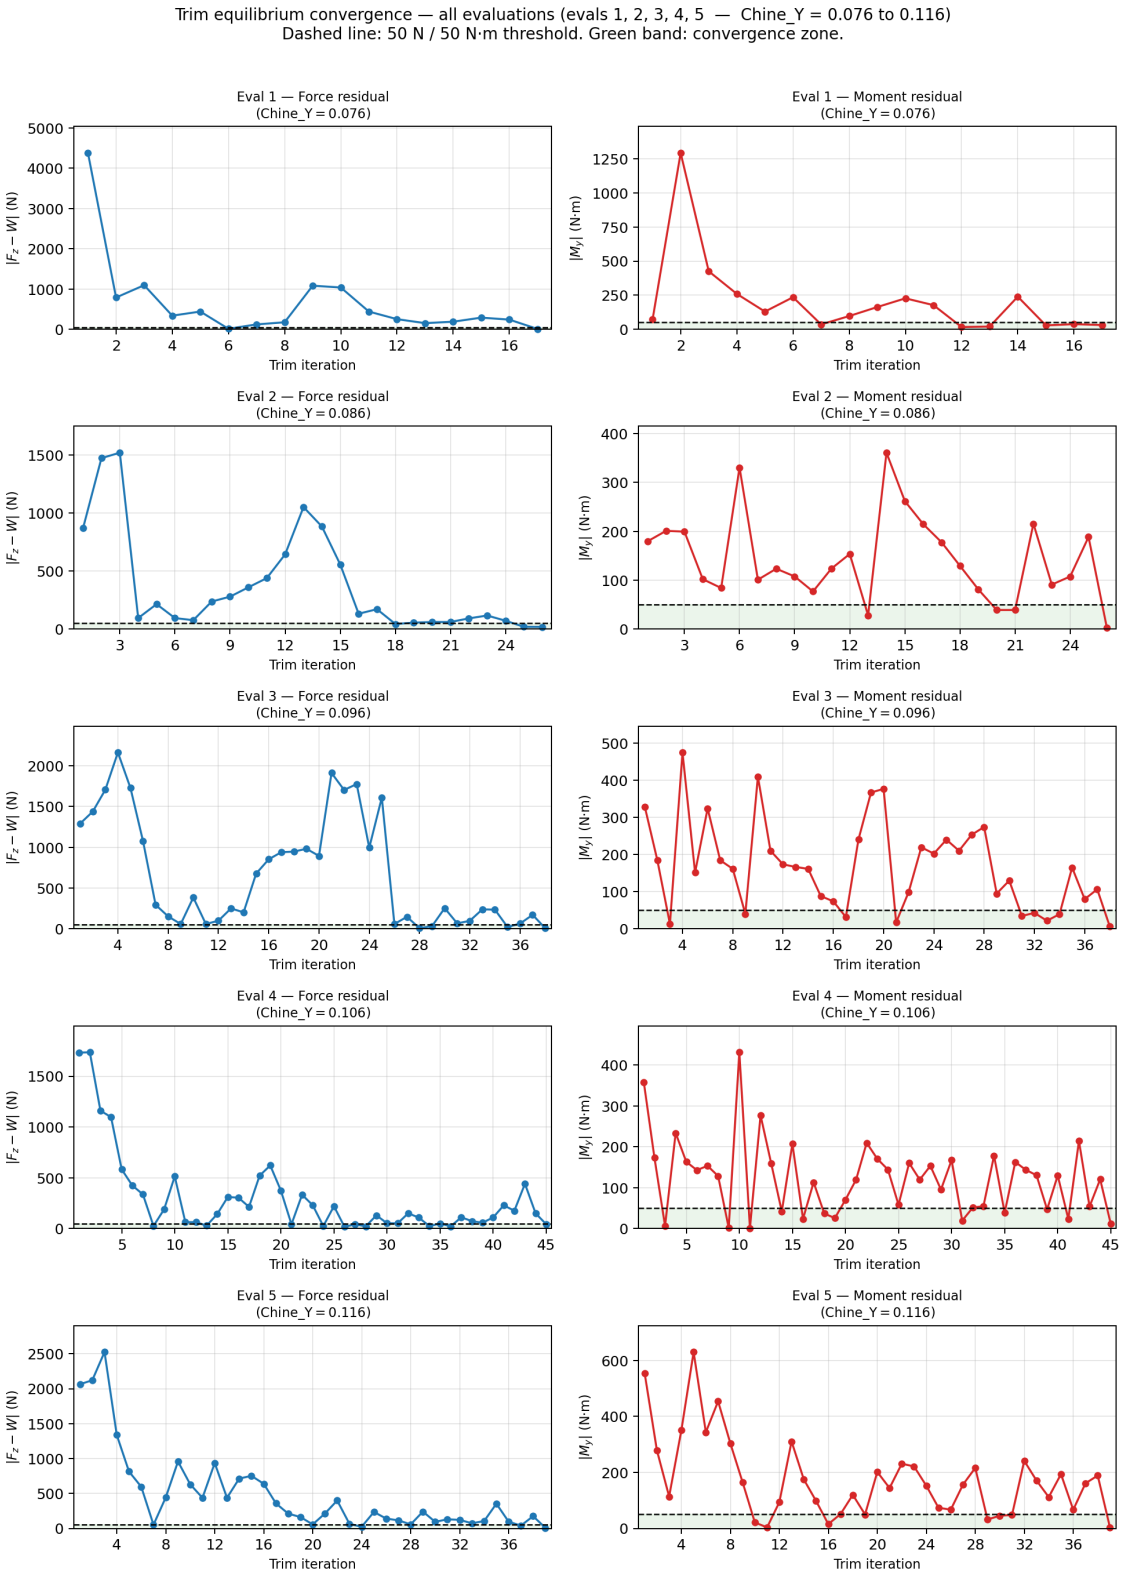


Figure C.2: Trim equilibrium convergence — evaluations 1, 2, 3, 4, 5 (Chine_Y = 0.076 to 0.116). Same layout as Figure C.1.

Appendix D

CFD Solver Convergence Histories

Figure D.1 shows the drag convergence history for the best trim iteration of each of the 10 optimization evaluations. Each panel corresponds to one hull geometry (Chine_Y value indicated). The Fluent solver reaches a stable drag plateau in all cases, confirming that the recorded drag values are not affected by insufficient CFD iteration. The dashed line in each panel marks the converged drag reported in Table 6.1. The consistent convergence behaviour across evaluations with very different hull geometries validates the robustness of the CFD setup and boundary conditions.

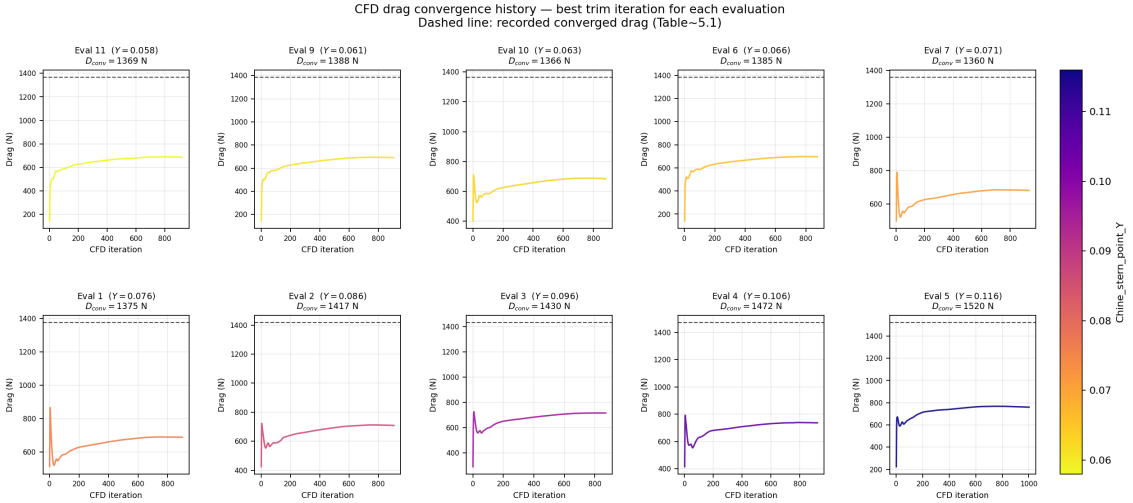


Figure D.1: CFD drag convergence history for the last trim iteration of each of the 10 optimization evaluations (eval 8 excluded). Each panel shows one hull (Chine_Y value indicated); the dashed line marks the recorded converged drag (Table 6.1). The colour scale maps Chine_Y from narrow chine (dark, purple) to wide chine (light, yellow). Note: for evaluations 5, 6, and 9, which were re-run to the 50 N criterion on a separate workstation, the figure shows the last locally available trim iteration; the converged drag values in Table 6.1 were taken from the completed re-run.

Appendix E

Full Mesh Sensitivity Table

Table E.1 extends Table 5.1 with absolute dimensionless coefficients. All coefficients use the Froude normalisation: $C_D = D/W$, $C_L = F_z/W$, $C_M = M_y/(W\,L_{wl})$, where $W = mg = 9810\,\mathrm{N}$ and $L_{wl} = 7\,\mathrm{m}$. Results are for the base hull at fixed trim ($\alpha = 0°$, $z = 0$) and are not directly comparable to the optimisation results (which use equilibrium trim). The percentage deviations are relative to the finest mesh (1.5 cm).

Table E.1: Full mesh sensitivity table: absolute dimensionless coefficients and percentage deviation from the finest mesh (1.5 cm reference). The selected mesh ($\Delta x = 3\,\mathrm{cm}$) is highlighted. Base hull at fixed trim ($\alpha = 0°$).

Δx (cm)	Elements	$C_D \times 10^{-2}$	ΔC_D (%)	$C_L \times 10^{-2}$	ΔC_L (%)	$C_M \times 10^{-2}$	ΔC_M (%)	Time (s)
1.5	2 229 372	6.815	ref	54.80	ref	−23.69	ref	13807
2.0	1 506 134	6.959	+2.11	56.09	+2.35	−24.64	+4.01	8192
2.5	1 329 728	6.932	+1.72	54.85	+0.09	−23.78	+0.38	7075
3.0	**967 971**	**7.000**	**+2.71**	**55.54**	+1.35	**−24.32**	+2.66	**7529**
3.5	874 897	7.181	+5.37	55.31	+0.93	−24.18	+2.06	3889
4.0	812 985	7.136	+4.71	55.30	+0.91	−24.30	+2.57	4184

Appendix F

Grasshopper Hull Parameters

Table F.1 lists the 29 geometric parameters (p_0–p_{28}) exposed as Grasshopper sliders in the parametric hull model. Each parameter controls one aspect of the hull shape. Parameter p_8 (`Chine_stern_point_Y`) is the variable optimised in this work; the remaining parameters are held at their base values throughout the optimisation campaign. In a future multi-parameter campaign, any subset of these sliders could be released as free variables.

Table F.1: Grasshopper hull parameters (p_0–p_{28}). The optimised parameter (p_8) is highlighted in bold.

Index	Parameter name
p_0	`Keel_run_start_Xrel`
p_1	`Keel_bow_X`
p_2	`Keel_depth_Z`
p_3	`Keel_middle_x`
p_4	`Transom_depth_Z`
p_5	`Bow_height_Z`
p_6	`Stem_angle`
p_7	`Chine_stern_bow_point_Z`
p_8	**`Chine_stern_point_Y`**
p_9	`Chine_side_exit_angle`
p_{10}	`Chine_side_entrance_angle`
p_{11}	`Chine_stern_Width`
p_{12}	`Side_bulkness`
p_{13}	`Rad_size`
p_{14}	`Rad_min_Z`
p_{15}	`Rad_max_Z`
p_{16}	`Radius_X_start`
p_{17}	`PB_Pt1_X`
p_{18}	`PB_Pt1_Y`
p_{19}	`PB_Pt1_Z`
p_{20}	`PB_Pt2_X`
p_{21}	`PB_Pt2_Y`
p_{22}	`PB_Pt2_Z`
p_{23}	`C1_Y_middle_point_X_position`
p_{24}	`C1_Y`
p_{25}	`Stern_point_Y_direction`
p_{26}	`C2_middle_point_X_direction`
p_{27}	`C2_middle_point_Y_direction`
p_{28}	`C2_middle_point_Z_direction`

Appendix G

Velocity Streamlines — Base Hull (Evaluation 4)

Figure G.1 shows the velocity streamline field for the base hull configuration (evaluation 4, Chine_Y = 0.106), viewed from above. The velocity magnitude reaches 7.04 m/s ($\approx 1.07\,V$) locally near the bow shoulder, while the turbulent wake is clearly visible as a region of reduced velocity behind the stern. The diverging wake pattern and the bow shoulder acceleration are characteristic features of the semi-planing regime at $Fr = 0.796$.

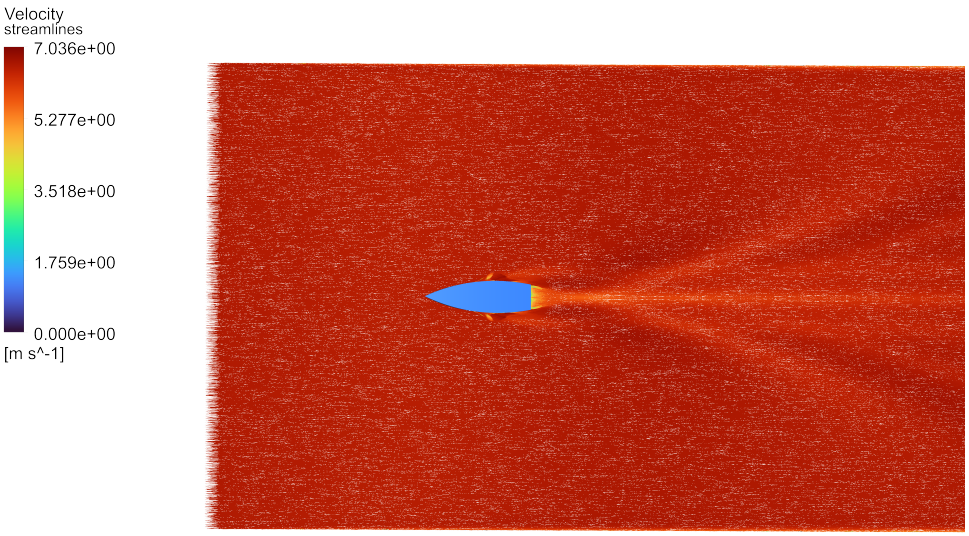


Figure G.1: Velocity streamlines on the free surface for the base hull (evaluation 4, Chine_Y = 0.106). Colour scale: velocity magnitude in $\mathrm{m\,s^{-1}}$.

Appendix H

Declaration of Use of Artificial Intelligence Tools

In accordance with the EPL guidelines for the responsible use of generative artificial intelligence, the following declaration is provided regarding the use of AI tools in the preparation of this thesis.

Writing assistance

An AI assistant was used as an aid in drafting and revising this manuscript, including assistance with phrasing, structure, and writing fluency in English. The author is responsible for all technical content, interpretations, and conclusions presented in this thesis.

Code development

An AI assistant was used as an aid during the development of the Python scripts implementing the automation and optimization pipeline (Chapter 4). The author is responsible for the overall workflow design and the validation of all results against the simulations.

General statement

The author affirms that this use of AI tools has been conducted in accordance with the principles of responsibility, authenticity, and transparency established by the EPL guidelines.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**