Optimization and Engineering
https://doi.org/10.1007/s11081-025-10055-4

RESEARCH ARTICLE

Development of an adjoint system-based hydrofoil optimisation framework using algorithmic differentiation

Rafael Tannenberg[1] · Karsten Hochkirch[2] · Andrea Walther[3] · Stephen R. Turnock[1] · Stephen W. Boyd[1]

Received: 21 May 2024 / Revised: 15 September 2025 / Accepted: 27 October 2025

Abstract

Various optimisation problems concern a component of a system, whose influence is so large that it significantly affects the state of the system. In these cases, an isolated optimisation of the component does not account for the changes in system state during the optimisation. This introduces inaccuracies. At the same time, the large influence of the component results in large potential performance gains. This requires detailed optimisation. The logical consequence is to model the whole system at every step of the optimisation and to use a large number of design variables. Both of these aspects can however increase computational time so significantly that the approach becomes infeasible. One such problem are hydrofoils in yacht racing. Hydrofoils are the equivalent to airfoils but operated underwater to lift the hull of a yacht out of the water. The design of the hydrofoils has an immense influence on the performance, the state and the trim (i.e. control) of the "yacht" system. To model this whole system, a stationary physics model of the entire yacht is developed. The model is integrated into a detailed optimisation routine that requires 70 design variables, which makes it prohibitively expensive to solve with derivative free methods. Therefore, a gradient-based optimisation strategy is developed, where the gradient is computed using the adjoint method. The adjoint method allows to compute the gradient independent of the number of input variables at a small cost. The adjoint method is only applied to the bottleneck of the yacht model using the algorithmic differentiation tool ADOL-C. The remainder of the model is differentiated using finite differences. The overall gradients are provided to the optimisation algorithm IPOPT. The optimisation strategy is applied to the AC75 America´s Cup class and used to optimise its hydrofoil for velocity made good (V_{MG}) in an upwind condition. The optimised foil shows significant improvement over the baseline foil and demonstrates the immense capabilities of adjoint system-based optimisation. Due to the vast efficiency of the adjoint method, the framework can be extended to optimise thousands of design variables.

Extended author information available on the last page of the article

Published online: 23 January 2026

Keywords Hydrofoil · Optimisation · Adjoint method · Algorithmic differentiation · System-based

Nomenclature

A_i	Area of segment i [m^2]
C_L	Lift coefficient
c_{ro}	Root chord [m]
D	Drag force [N]
$\boldsymbol{dv}$	Hydrofoil design variables [m, °]
F	VPP Function [kts]
$\boldsymbol{F_i}$	Force at control point i [N]
F	Function
F_i	Function i
f_x	Force in x-direction [N]
f_y	Force in y-direction [N]
f_z	Force in z-direction [N]
$\boldsymbol{fm}$	force- and moment vector [N, Nm]
L	Lift force [N]
l_i	Lagrangian multiplier i
$\boldsymbol{l_i}$	Length of segment i [m]
l_{ext}	Extension lengths [m]
l_{spa}	Semi span [m]
$N_{\mathrm{i,p}}(u)$	Nonrational B-spline basis function in *u*-direction
$N_{\mathrm{j,q}}(v)$	Nonrational B-spline basis function *v*-direction
m_x	Moment around x-axis [Nm]
m_y	Moment around y-axis [Nm]
m_z	Moment around z-axis [Nm]
p, q	Degrees of NURBS surfaces in *u*- and *v*-direction
$\boldsymbol{p}$	Point a function is evaluated for
$\boldsymbol{P}_{\mathrm{i,j}}$	Control points of NURBS-Surfaces [m]
r_{tap}	Taper ratio
$\boldsymbol{r_{i_0 j}}$	Vector from node i_0 to control j
$\boldsymbol{r_{i_1 j}}$	Vector from node i_1 to control j
$r_{i_0 j}$,$r_{i_1 j}$	Magnitudes of $\boldsymbol{r_{i_0 j}}$and$\boldsymbol{r_{i_1 j}}$
$\boldsymbol{S}$	NURBS surface [m]
$\boldsymbol{u_\infty}$	Unit vector in the direction of the freestream
u, v	Directions of NURBS surfaces
$\boldsymbol{V_i}$	Local velocity at control point i [m s^{-1}]
V_{MG}	Velocity made good [kts]
V_S	Boat speed [kts]
V_{TW}	True wind speed [kts]
$\boldsymbol{V_\infty}$	Free-stream velocity [m s^{-1}]
$\boldsymbol{v_{ji}}$	Influence of horseshoe vortex j on control point i [m^{-1}]
$w_{\mathrm{i,j}}$	Weights of control points of NURBS surfaces
x_i	Input variable i
β_t	True wind angle [°]

Γ_j	Strength of horseshoe vortex j [$m^2\ s^{-1}$]
δ_e	Elevator rake [°]
δ_{ra}	Rake angle [°]
δ_r	Rudder angle [°]
δ_{win}	Wing rake angle [°]
λ	Leeway angle [°]
$\boldsymbol{\xi}$	Free variable vector
ρ	Density [kg m^{-3}]
τ	Flat parameter
as	Automatically spaced
AVL	Athena Vortex Lattice
BFGS	Broyden–Fletcher–Goldfarb–Shanno
CFD	Computational Fluid Dynamics
CoG	Centre of Gravity
DOF	Degree of freedom
ETNZ	Emirates Team New Zealand
FEA	Finite Element Analysis
IMS	International Measurement System
IPOPT	Interior Point Optimizer
LLT	Lifting line theory/method
LRPP	Luna Rossa Prada Pirelli
NURBS	Non-uniform rational B-Splines
RANS	Reynolds-averaged Navier Stokes
VPP	Velocity prediction program

1 Introduction

Various optimisation problems concern a component of a system, whose influence is so large that it significantly affects the state of the system. In these cases, an isolated optimisation of the component does not account for the changes in system state during the optimisation. This introduces inaccuracies and uncertainty but only requires a single evaluation of the component model per iteration of the optimisation. Modelling the whole system mitigates the inaccuracies but requires significantly more computational resources as the whole state of the system must be computed at every iteration during the optimisation process. Solving for the state requires numerous component model evaluations itself plus the modelling of the other components in the system. Simultaneously, the large influence of the component results in large potential performance gains, which calls for detailed optimisation. To fully exploit the design space a high number of design variables is required. However, the run time of an optimisation often scales with the number of variables, which makes detailed design prohibitively expensive. This is especially the case when the whole system is modelled.

One such component are the hydrofoils on racing yachts. Hydrofoils are the equivalent to airfoils but operated underwater to lift the hull of a yacht out of the water. Their effective lift/drag-ratio allows the yacht to reach speeds much higher than with

the hull in the water. The design of the hydrofoils thereby has an immense influence on the performance, the state and the trim (i.e. control) of the system, i.e. "the yacht". This was shown by Paulin et al. (2015) at the example of a C-Class catamaran. Two promising hydrofoil designs were compared using a stationary physics model of the yacht. The model predicted differences in boat speed of up to 20%. Other state variables, such as the leeway angle, changed up to 200%. The control-variable rake changed up to 220%. Conventional hydrofoil optimisations are performed in isolation of the yacht system neglecting these immense changes in yacht state. They are usually based on the state of an initial hydrofoil design. Hence, new candidates are evaluated for a state that they will never experience in reality.

In addition, the high cost of the hydrodynamic model of the foil and/or the optimisation strategy limits the number of design variables. Examples are the optimisations performed by Meneghello et al. (2016) and Ploe (2018). Meneghello used a lower fidelity vortex lattice method, while Ploe modelled the hydrofoil forces with high fidelity computational fluid dynamics (CFD). Despite both contributions employing surrogate models, their approaches were limited to seven and two design variables, respectively. A detailed optimisation of the spine- and planform of the foil, as well as its twist distribution, requires a much higher number of design parameters (70 – several hundreds if section design is to be included). The only feasible option for more than 70 design variables is a gradient-based algorithm combined with the adjoint method for gradient computation. The adjoint method computes the gradient independent of the number of input variables at a small cost but requires extensive implementation efforts. Adjoint optimisations have been performed in many fields including Biology (Hovland et al. 1997), Medicine (Jee et al. 2005), Physics (Kim et al. 2006) and Engineering. Examples in engineering range from electrical engineering (Hart et al. 2006), over controller design (Röbenack 2007), to structural optimisation (Tadjouddine et al. 2006). In the context of hydrofoil optimisation, Garg et al. (2017) conducted an adjoint hydrostructural optimisation using high fidelity CFD and finite element analysis (FEA). The use of the adjoint method resulted in the ability to optimise 200 design variables. This optimisation was, however, also carried out in isolation of the yacht system.

This work presents the development of an adjoint hydrofoil optimisation strategy where the whole yacht system is considered. The system is modelled with a six-degree-of-freedom stationary physics model of the yacht developed in the velocity prediction program (VPP) FS-Equilibrium. The adjoint version of the VPP model is derived using the algorithmic differentiation tool ADOL-C (Walther and Griewank 2020). The adjoint method is however only applied to the bottleneck of the VPP, which corresponds to the parametric model of the hydrofoil and the lifting line method computing the hydrofoil forces. This represents the first part of the gradient. The remainder of the program (solver and other force modules) is differentiated using the simpler finite difference method (second part of the gradient). This saves implementation time and preserves modularity. The two gradient parts are combined using the chain rule and are provided to the optimisation algorithm IPOPT (Interior Point Optimizer) (Wächter and Biegler 2006) at every step of the optimisation.

The methodology is applied to an AC75 type yacht. The AC75 is the foiling monohull class currently sailed in the America´s Cup, the Formula 1 of sailing. Figure 1

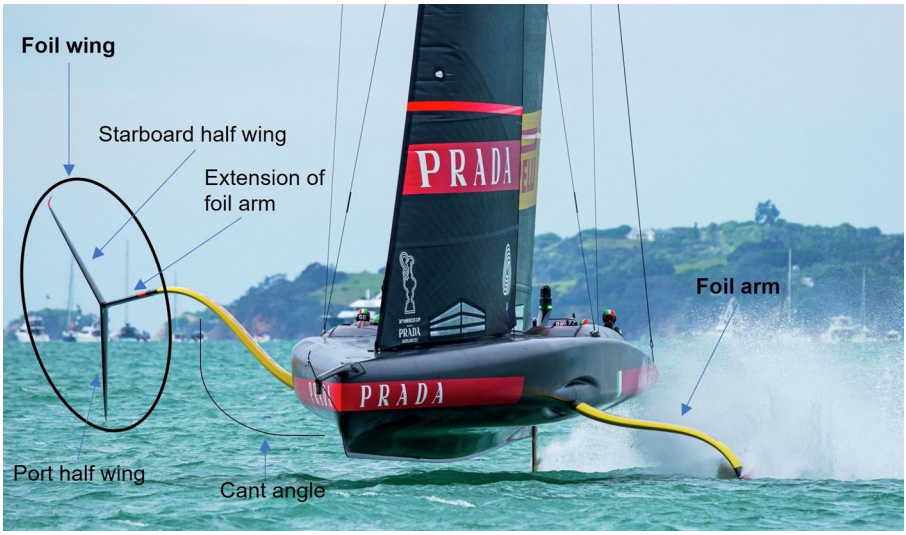

Fig. 1 AC75 Luna Rossa based on Gattini (2020)

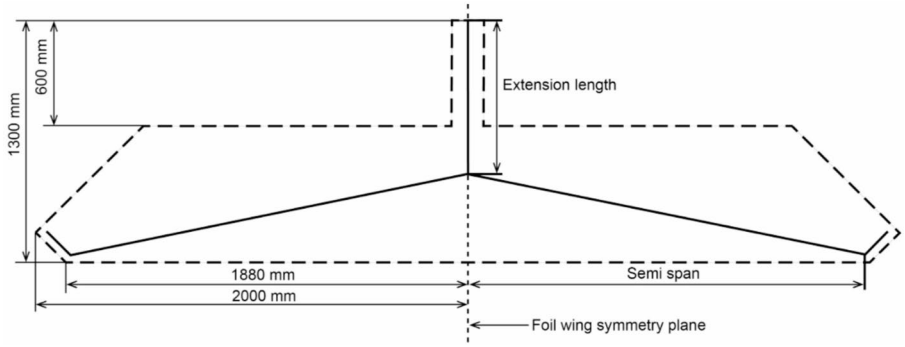

Fig. 2 Foil box

shows an AC75 operating with the windward foil out of the water to provide righting moment and the leeward foil in the water to produce hydrodynamic lift and side force. The main foils consist of two parts which are the foil arm and the foil wing. The design of the foil wing is open to development inside the `Foil box´ (Fig. 2) and must be symmetric around the `Foil wing symmetry plane´. Furthermore, the foil wing has to weigh exactly 921 kg. The wings are built from a combination of steel and lead to accommodate this weight while having a low volume and wetted surface area. Some of the weight can be stored in a bulb and one or two flaps are used for force control. In the last America's Cup, Emirates Team New Zealand (ETNZ) and American Magic used bulbs and wings with small chord while Luna Rossa Prada Pirelli (LRPP) and INEOS Team UK opted wings with no bulb and larger chord. ETNZ´s foil had no anhedral angle (T-Foil) and a single flap, the other teams chose Y-Foils with two flaps.

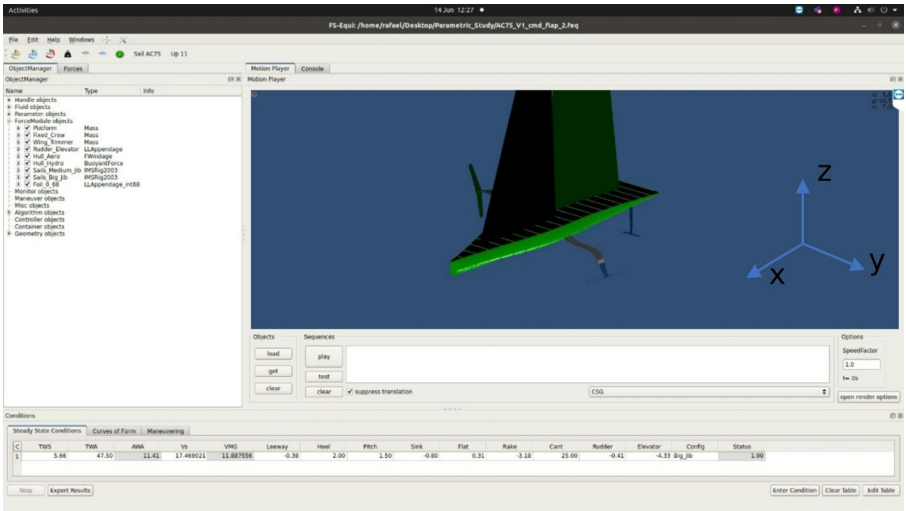

Fig. 3 FS-Equilibrium´s graphical user interface showing a generic AC75

The developed framework is first used to optimise five design variables describing the general design of the AC75 hydrofoil. The results are compared to the findings of a parametric study recently published by Tannenberg et al. (2023) which acts as a first test of the routine. The parametric study used the same yacht model and variables. As a second test, the twist-distribution of a straight T-Foil with no taper ratio is optimised. This case is analysed to determine if the optimiser can tailor the twist-distribution such that the desired elliptical lift-distribution is achieved. With the two tests completed, a detailed optimisation of the entire foil is performed. The detailed optimisation has 68 design variables describing the spine- and the planform of the foil as well as the twist distribution in detail.

2 Velocity prediction model

The physics model of the AC75 is developed in the VPP FS-Equilibrium (Fig. 2) and is briefly introduced in this section, for a more thorough description see Tannenberg et al. (2023). For more information on FS-Equilibrium see Hochkirch (2018). FS-Equilibrium is a workbench for sailing yacht performance prediction and can compute stationary and dynamic sailing states. “Force modules” are used to model the different components of the yacht and return the forces and moments they produce. Once all forces and moments equate to zero in the six degrees of freedom (DOF), a steady sailing state is reached. This steady sailing state is determined using a Newton-Raphson method. The method modifies state and trim variables until force equilibrium is found. FS-Equilibrium has been applied to hydrofoiling yachts before such as the C-Class Catamaran Groupama (Paulin et al. 2015), the International Moth (Eggert 2018), the AC50 (Hansen et al. 2019) and the QFX Lake Racer (Melis et al. 2022).

FS-Equilibrium is used to compute the steady sailing states of the AC75, by balancing the six forces and moments

$$fm = (f_x, f_y, f_z, m_x, m_y, m_z)^T \tag{1}$$

Table 1 shows the state and trim variables used to balance the forces of the AC75 model. These variables do not exclusively influence one degree of freedom but are the variables with the highest influence on the respective degree of freedom. For simplification reasons the foil wings are raked rather than controlled with flaps as on the real yachts. The remaining state variables, heel angle, pitch angle and ride height (sink) are fixed for increased robustness.

The problem is expressed with Lagrange multipliers ($l_i \neq 0$) in the function

$$F = -V_S + l_0 \sum f_x + l_1 \sum f_y + l_2 \sum f_z + l_3 \sum m_x + l_4 \sum m_y + l_5 \sum m_z\,, \tag{2}$$

where the first order condition for the optimum is

$$F \overset{!}{=} \min \rightarrow \quad \nabla F \overset{!}{=} 0. \tag{3}$$

This leads to a non-linear system of equations:

$$\begin{aligned}
&F_0 = \frac{\partial F}{\partial V_S} \overset{!}{=} 0, \quad F_1 = \frac{\partial F}{\partial \lambda} \overset{!}{=} 0, \quad F_2 = \frac{\partial F}{\partial \delta_{\mathrm{ra}}} \overset{!}{=} 0, \\
&F_3 = \frac{\partial F}{\partial \tau} \overset{!}{=} 0, \quad F_4 = \frac{\partial F}{\partial \delta_e} \overset{!}{=} 0, \quad F_5 = \frac{\partial F}{\partial \delta_r} \overset{!}{=} 0, \\
&F_6 = \frac{\partial F}{\partial l_0} \overset{!}{=} 0, \quad F_7 = \frac{\partial F}{\partial l_1} \overset{!}{=} 0, \quad F_8 = \frac{\partial F}{\partial l_2} \overset{!}{=} 0, \\
&F_9 = \frac{\partial F}{\partial l_3} \overset{!}{=} 0, \quad F_{10} = \frac{\partial F}{\partial l_4} \overset{!}{=} 0, \quad F_{11} = \frac{\partial F}{\partial l_5} \overset{!}{=} 0.
\end{aligned} \tag{4}$$

This non-linear system of equations is solved using the Newton-Raphson method. This solver set-up allows the direct inclusion of a further optimisation variable. This feature is not yet used but can be relevant for including the optimisation of an additional state variables in the future. For the current problem, the Newton-Raphson could also be directly employed for force balancing. This is more computationally efficient.

Table 1 Solver set-up of the AC75 physics model

DOF	Balancing variable with the highest influence on the DOF
f_x	boat speed V_S (state)
f_y	leeway angle λ (state)
f_z	rake angle δ_{ra} (trim)
m_x	sail-power (flat) τ (trim)
m_y	elevator rake δ_e (trim)
m_z	rudder angle δ_r (trim)

The force modules of the AC75 are presented in Table 2. For details see Tannenberg et al. (2023). The forces and moments generated by the hydrofoils are modelled using an advanced lifting method (LLT) as presented in the next section.

3 Hydrodynamic foil model

The forces and moments produced by the main hydrofoils of the AC75 are computed with a lifting line method. The method was originally derived by Prandtl (1918) and is a non-viscous method for 3D foil force prediction. Viscosity is of importance in stalled cases or when the flow separates. For the desired operation of a hydrofoil, this is not the case, and the lifting line method has an accuracy similar to viscous methods such as Reynolds-averaged Navier Stokes (RANS). RANS methods require hours of computational time on multiple cores for a single foil force evaluation and are not suited for the use within yacht models that involve numerous such evaluations to compute the sailing state. Lifting line methods on the other hand solve in milliseconds on a single core and are the most common method for foil force prediction in stationary and dynamic yacht simulators (e.g. SumToZero 2021). The lifting line method uses vortices in the free stream to represent how a foil section affects its surrounding flow. Multiple vortices are oriented along a line that follows the span of the foil. For a comprehensive introduction to lifting line methods see Katz and Plotkin (1991). A variation of Prandtl´s method is used for this study. This advanced method was developed by Phillips and Snyder (2000) and employs horseshoe vortices at every station along the span and a vectorised form of the Kutta-Jukowski-Law. Each horseshoe vortex is formed of a bound vortex and two semi-infinite vortices. The vectorised lifting law includes the influence of the bound vortices on each other. This is required to model the effects of leeway, sweep and anhedral. It also includes the influence of multiple surfaces on each other. The method solves for the strength of each horseshoe vortex Γ_j by relating two definitions of the force generated at every segment. The force produced by each bound vortex is computed with

$$d\boldsymbol{F_i} = \rho \Gamma_i \boldsymbol{V_i} \times d\boldsymbol{l_i}\,, \tag{5}$$

where the local velocity at every control point is calculated from

Table 2 AC75 force modules

Component	Model	Comments
Platform gravity	Gravity model	Weight and CoG from AC75 rule
Crew gravity	Gravity model	Weight from rule, CoG estimated
Hull aero	Coefficient-based	Coefficients from Hoerner (1965)
Sails aero	Tuned IMS2003	Factor of 1.3 on driving force
Rudder hydro	Lifting Line	Based on Phillips and Snyder (2000)
Main foil hydro	Lifting Line (adjoint)	Same as for Rudder, but adjoint

$$\boldsymbol{V_i} = \boldsymbol{V_\infty} + \sum_{j=1}^{N} \Gamma_j \, \boldsymbol{v_{ji}} \, . \tag{6}$$

The velocity $\boldsymbol{v_{ji}}$ is the induced velocity of horseshoe vortex j at control point i normalised by the vortex strength and is computed with

$$\boldsymbol{v_{ji}} = \frac{1}{4\pi} \left[\frac{\boldsymbol{u_\infty} \times \boldsymbol{r_{i_1 j}}}{r_{i_1 j}(r_{i_1 j} - \boldsymbol{u_\infty} * \boldsymbol{r_{i_1 j}})} + \frac{(r_{i_0 j} + r_{i_1 j})(\boldsymbol{r_{i_0 j}} \times \boldsymbol{r_{i_1 j}})}{r_{i_0 j} r_{i_1 j}(r_{i_0 j} r_{i_1 j} + \boldsymbol{r_{i_0 j}} * \boldsymbol{r_{i_1 j}})} - \frac{\boldsymbol{u_\infty} \times \boldsymbol{r_{i_0 j}}}{r_{i_0 j}\,(r_{i_0 j} - \boldsymbol{u_\infty} * \boldsymbol{r_{i_0 j}})} \right] \tag{7}$$

where the first term in the brackets is the outbound semi-infinite vortex influence, the second term the bound vortex influence and the third term the inbound semi-infinite vortex influence of segment i on colocation point j. The vectors $\boldsymbol{r_{i_0 j}}$ and $\boldsymbol{r_{i_1 j}}$ point from the semi-infinite vortices to the colocation point and $r_{i_0 j}$ and $r_{i_1 j}$ are the lengths of these vectors. The bound vortex term is excluded when $i = j$, as the bound vortex is not causing downwash on its own segment. Relating the lift based on the sectional lift coefficient to the vectorised vortex lifting law results in a non-linear equation system. This is solved iteratively but is potentially unstable. Hence, the method was reduced to a linear problem through a linear lift-curve slope and yields the following system of equations

$$\rho \, \Gamma_i \left| \left(\boldsymbol{V_\infty} + \sum_{j=1}^{N} \Gamma_j \, \boldsymbol{v_{ji}} \right) \times \mathrm{d}\boldsymbol{l_i} \right| - \frac{1}{2} \rho \, \boldsymbol{V_\infty}^2 \, C_L \, (\boldsymbol{V_i}) \, dA_i = 0. \tag{8}$$

For a more general introduction to hydrofoils and associated models see Molland and Turnock (2022). The two half wings of the foil are discretized with 30 stations each, while the extension uses 10, all with equal spacing. This is in contrast to the 40 stations per part with cosine-spacing proposed by Phillips and Snyder (2000), but delivers much more robust results in the optimisation. The cosine-distribution has led to unrealistic optimal geometries especially in the region of the junction. This was amplified by the small spacing at the junction caused by the high number of stations. Parts of the foil that are above the waterline are not considered (Fig. 4). The loss in lift caused by the presence of the free surface is modelled with a second transformed foil. The biplane analogy (Faltinsen 2005) is used, which is best suited for the design space and conditions explored.

The force module of the main hydrofoil additionally computes the forces and moments from the weight of the foil based on an integration of the sectional areas along the span using Simpson´s rule.

An average density of 8305 kg m^{-3} of the involved materials (75% high density steel, 20% lead and 5% hydraulic oil to represent the control systems) is used. A bulb is automatically sized to meet the weight requirement of 921 kg. The centre of gravity of the foil is determined by taking moments of the sectional areas around the junction of the foil arm and the foil arm extension. The resistance of the potential bulb is computed using a form factor approach. The lifting line method combined with the form factor captures the general trends well, but loses accuracy in the region of the

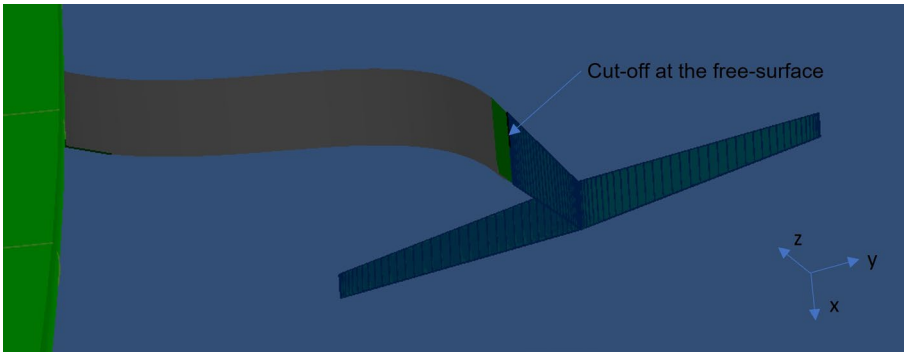

Fig. 4 AC75 foil discretised with 30 stations per half wing and 10 stations at the extension using an equal distribution

bulb. Implementing a slender body model would enhance the accuracy in the area, but should also be treated with caution and results should be validated with viscous CFD simulations. Furthermore, wave-making drag and spray drag are not/not yet considered. The model is available in normal mode for finding the steady sailing state and in adjoint mode for gradient computation as explained later.

4 Parametric model

The parametric model serves to generate the shape of the hydrofoil from the set of design variables. The lifting line method reads the geometrical information of the hydrofoil from ruled surfaces which entail for example the position of each segment, its chord length and twist. The ruled surfaces are provided from the parametric model in the form of NURBS (Non-rational uniform B-spline) surfaces. NURBS surfaces are defined as

$$S(u,v) = \frac{\sum_{i=0}^{n} \sum_{j=0}^{m} N_{\mathrm{i,p}}(u)\, N_{\mathrm{j,q}}(v)\, w_{\mathrm{i,j}} \boldsymbol{P}_{\mathrm{i,j}}}{\sum_{i=0}^{n} \sum_{j=0}^{m} N_{\mathrm{i,p}}(u)\, N_{\mathrm{j,q}}(v)\, w_{\mathrm{i,j}}} \qquad for \;\; 0 \leq u,v \leq 1. \quad (9)$$

where $N_{\mathrm{i,p}}(u)$, $N_{\mathrm{j,q}}(v)$ are the nonrational B-spline basis functions, $\boldsymbol{P}_{\mathrm{i,j}}$ are the control points and $w_{\mathrm{i,j}}$ are the weights of the control points. The directions u and v are alongside the surface (u spanwise, v chordwise) and p and q are the degrees in the u- and v-direction, respectively. See Piegl and Tiller (1997) for more detail. The cross-sections of the foil at every station (i.e. the "thickness" of the foil) are supplied through their sectional properties (i.e. lift curve slopes, the zero lift angles and the profile drag coefficients). This is discussed in detail in Tannenberg et al. (2023).

Figure 5(a) shows the simple parametric model which was used in the parametric study. It describes the hydrofoil with the parameters extension length l_{ext}, semi span l_{spa}, root chord c_{ro}, taper ratio r_{tap} and wing rake δ_{win}. This simple model is used in a first optimisation to compare the optimisation results with the findings of the parametric study.

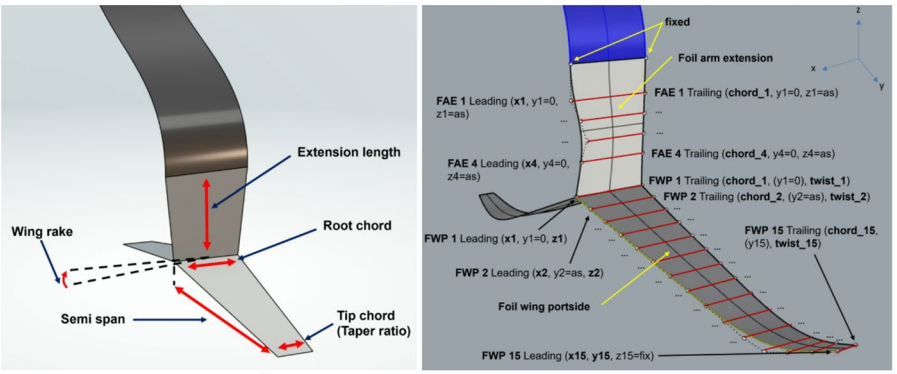

Fig. 5 Simple parametric model with 5 variables (**a**), detailed model with 68 variables (**b**); design variables bold; as = automatically spaced

The twist-distribution optimisation and the detailed optimisation of the hydrofoil are performed using the advanced parametric model shown in Fig. 5(b). This model is based on 68 design variables and allows thorough optimisation of the spine and plan form of the foil as well as the twist distribution. The extension of the foil arm has four stations which are defined by the x-coordinates of every leading edge point and the chords. The y-coordinates and the twists are set to zero due to the symmetry requirement. The stations are equally spaced in z-direction between the end of the foil arm and the root of the foil wing. This results in 8 variables for the extension. The portside half wing has 15 stations. Every station is defined by the x-and z-values of the leading edge, a chord and a twist variable. Station 15 additionally has a y-variable that defines the span of the wing. Its z-value is however fixed at the bottom of the foil box. The y-value of station 1 is zero due to the symmetry constraint. The intermediate stations are spaced equally in between. This results in 60 variables for the portside half wing, which is mirrored around the foil wing symmetry plane to generate the starboard half wing. This model is sophisticated enough to re-create the shapes of the hydrofoils seen in the 36th America's Cup as shown in Fig. 6 representing the ETNZ (a) and LRPP (b) hydrofoil, respectively. The models are based on the C++library TinyNURBS (Jayaraman 2022) and use the python library NURBS-python (Bingol 2022) for visualisations. The parametric models are internally coupled to the lifting line method to allow seamless gradient computation as discussed in the next section.

5 Model validation

The correct functioning and accuracy of the hydrodynamic model of the foil and the entire physics model of the yacht are crucial for reliable optimisation results. The hydrodynamic model has been validated in terms of lift/drag-ratio. The whole yacht model was validated with real life race data from the last America´s Cup (America´s Cup, 2020). Both showed good agreement as discussed in Tannenberg et al. (2023). Herein, a more thorough validation is performed for the lift-distribution computed by the hydrofoil model. The predicted lift-distributions are benchmarked against the

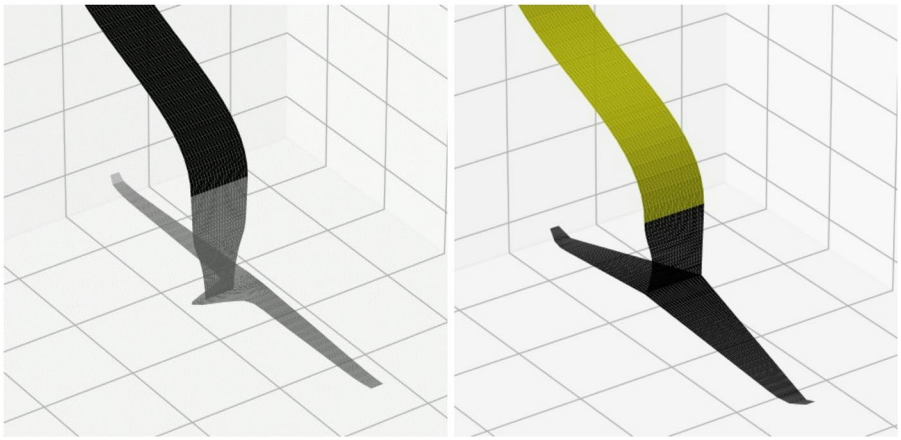

Fig. 6 Remodelled hydrofoils of ETNZ (**a**) and LRPP (**b**)

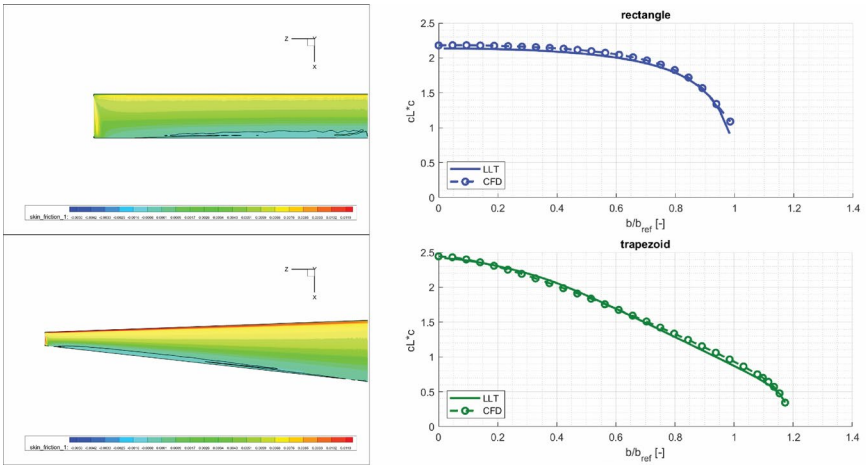

Fig. 7 Lift-distributions of a non-tapered and a tapered half wing (Hospodář et al. 2022)

distributions computed by Athena Vortex Lattice (AVL) (Drela and Youngren 2022) and those found in the literature for two standard foil shapes. It is also analysed if the computation of the weight and the centre of gravity of the foils is executed correctly.

The lift-distribution along the span of the foil is an important characteristic in hydrofoil design with respect to induced drag. Figure 7 shows the lift-distributions for a half-wing with no taper and one with a taper ratio of 0.25, both with constant twist. The lift-distributions are predicted by a lifting line method and a computational fluid dynamics code (Hospodář et al. 2022). The straight wing shows a lift-distribution with a more rectangular shape than the desired elliptical shape due to the constant chord. The taper of the second foil reduces the lift generated towards the tip which results in a more triangular shape.

Two similar foils are analysed with the lifting line method used in the optimisation framework and AVL for further comparison. The results are shown in Fig. 8. Both, the lifting line method and AVL predict the rectangular and the triangular lift-distributions observed in Fig. 7. Slight differences between the two can be seen towards the tip in the rectangular distribution and in the middle between the root and the tip of the tapered foil. The lift-distribution of the tapered foil predicted by the lifting line method also shows a slight drop in lift at the root. This is not predicted by the other models. In general however, the methods show good agreement, and the lifting line method is hence deemed suitable for lift-distribution prediction.

The hydrofoil force module also computes the weight and the centre of gravity of the foil. The sectional area of the foil is integrated along the span which results in a 2.5D approach rather than a full 3D approach. Moments of the sectional areas are used to compute the centre of gravity. This is perfectly acceptable if the spine of the foil is straight but can lose accuracy if the spine is curved. Therefore, foils with different amounts of curvature have been tested and compared to a full 3D approach implemented in a commercial computer aided design software. The weight error was 3.2% for the foil with the highest curvature and 0% for a straight foil. The centre of gravity error was 4% for the highest curvature foil and 0% for the straight foil. This is deemed sufficient.

6 Optimisation strategy

This section presents the logic of the VPP driven hydrofoil optimisation routine while the details of the gradient computation are discussed in the next chapter. Figure 9 shows the flow of the proposed routine. In a first step, the optimisation point is defined. This is the true wind angle (β_t) and true wind speed (V_{TW}) the yacht sails in and hence the condition for which the foil is optimised. At a later stage, multiple of those points will be considered to optimise for a broader range of conditions. At this stage however, only one condition can be assessed in one optimisation run. In the second step, the design parameters of an initial foil are defined. The parameters are passed to the parametric model which generates the initial hydrofoil geometry. The VPP is now run solving for the equilibrium sailing state of the AC75 in the set condition with the initial hydrofoil. This is an iterative process where the hydrofoil forces are calculated numerous times by the advanced lifting line method for chang-

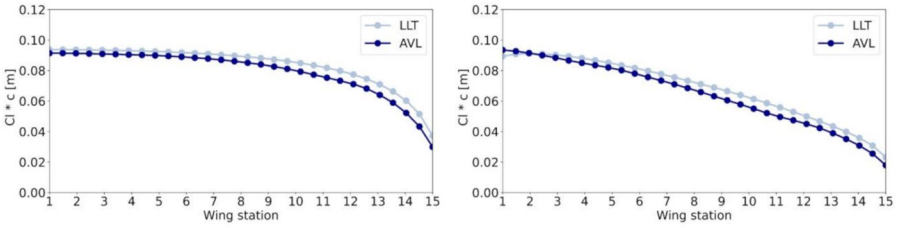

Fig. 8 Lift-distribution predicted by the lifting line method and AVL for the constant chord wing (**a**) and the tapered wing (**b**)

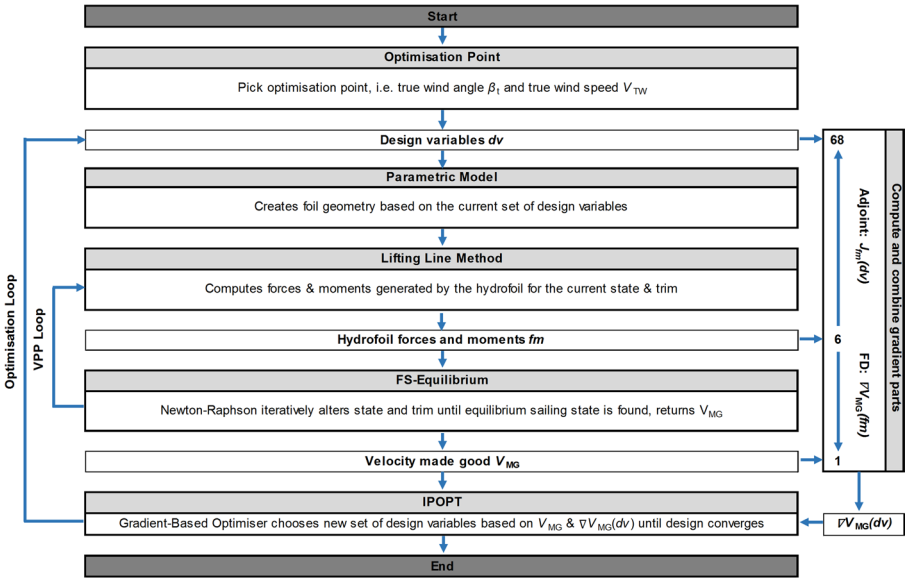

Fig. 9 Adjoint VPP-driven optimisation routine

ing states and trims (VPP loop). The lifting line method is directly coupled to the parametric model from which it receives the geometric information. Once a valid equilibrium sailing state is found, the VPP returns the achieved V_{MG} and passes it to the optimisation algorithm IPOPT. The optimisation is currently performed for V_{MG} but with a fixed β_t. This means optimising for V_{MG} is the same as optimising for V_{S}, but β_t will be included as an optimisation variable in the future. Additionally, the cant angle will be included as an optimisation variable. It is currently fixed to the position that generates the highest righting moment, while preventing the tip from piercing the surface.

IPOPT also requires the gradient of the function (VPP including lifting line and parametric model) which is computed next and passed to the optimiser. This three-step process is described in detail in the next section. Based on the V_{MG} and the gradient, IPOPT computes a set of design variables which results in a higher V_{MG}. This new set is fed back to the parametric model which updates the geometry and the VPP is run again. The new V_{MG} and gradient are fed to IPOPT again which comes up with an even better performing set of design variables (Optimisation loop). This process is repeated until the design has converged to the optimum. During this process, IPOPT also considers the bounds on the independent variables. These are treated as linear constraint functions. For the details on the inner workings of IPOPT see Wächter and Biegler (2006).

The most expensive part of this routine is the computation of the gradient. The runtime required for it increases with every design variable. This makes the VPP driven optimisation of a high number of variables infeasible when the gradient is approximated with finite differences. The next section presents an innovative approach that allows to compute the gradient at a small run time independent of the number of

input variables. As a result, thousands of variables can be optimised efficiently. Interior point methods also require second order derivatives. These are approximated by IPOPT using a BFGS (Broyden–Fletcher–Goldfarb–Shanno) method.

7 Gradient computation

The gradient describes the direction and rate of fastest increase of a scalar function ff at a point $\mathrm{p} = (\mathrm{x1}, \ldots, \mathrm{xn})\mathrm{p} = (\mathrm{x1}, \ldots, \mathrm{xn})$. For functions with multiple outputs, the gradients form the Jacobian matrix (see Eq. 13). A simple way to approximate gradients is finite difference: the function is evaluated for an initial input, then each design variable is perturbed once. For nn input variables and one output, $\mathrm{n} + 1\mathrm{n} + 1$ evaluations are required, which quickly becomes expensive (e.g. 69 runs for 68 variables). Alternatively, gradients can be obtained via forward differentiation using the chain rule, either manually or with automatic differentiation tools. Like finite difference, this requires one derivative per input variable and is independent of the number of outputs. For problems with many inputs and few outputs, reverse differentiation (the adjoint method) is preferable: the cost then depends only on the number of outputs. This makes it highly effective for optimisation with many design variables (Griewank and Walther 2008).

Differentiating a computer code however is a complex and time-consuming process even if algorithmic differentiation tools are used. Finite differences in contrast are simple to implement. To reduce development time, while making the run time independent of the number of input variables a hybrid approach was developed. The adjoint method is only applied to the bottleneck of the routine, while finite differences are used where it does not significantly affect run time. Consequently, the gradient parts are joined together. The resulting gradient computation runs at a small multiple of a single normal evaluation of the VPP model.

The bottleneck of the VPP driven routine is the parametric model and the lifting line method. This includes the routine which computes the bulb size and forces. Here the Jacobian has to be computed for the six forces and moments $\boldsymbol{fm}$ the hydrofoil produces with respect to the large number of design variables $\boldsymbol{dv}$. In case of the advanced parametric model, this results in 68 input variables and 6 output variables. Applying the adjoint method here means the Jacobian can be computed by evaluating six partial derivatives only. This holds true for any number of input variables. For the simple parametric model with five input variables $\boldsymbol{dv} = (l_{ext},\ l_{spa},\ c_{ro},\ r_{tap},\ \delta_{\ win})^T$ the Jacobian looks as follows

$$J_{\boldsymbol{fm}}(\boldsymbol{dv}) = \begin{bmatrix} \nabla {f_x}^T(\boldsymbol{dv}) \\ \nabla {f_y}^T(\boldsymbol{dv}) \\ \nabla {f_z}^T(\boldsymbol{dv}) \\ \nabla {m_x}^T(\boldsymbol{dv}) \\ \nabla {m_y}^T(\boldsymbol{dv}) \\ \nabla {m_z}^T(\boldsymbol{dv}) \end{bmatrix} \tag{10}$$

where the rows of $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ are the transposed gradients of every force/moment with respect to $\boldsymbol{dv}$. The Jacobian $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ of the big parametric model is computed in the same fashion but has 68 columns. The Jacobian is derived using the algorithmic differentiation tool ADOL-C. Thereby, the full source code must be available and has to be prepared for differentiation. This means all pre-compiled libraries have to be replaced in the first step before the whole routine is templated to facilitate a special datatype required for differentiation. This datatype allows the storage of primal values and to compute the derivatives. The templating includes any part of the code that influences the gradient, so from the overlaying lifting line class down to the functions of vector- and matrix libraries, in total 8000 lines of code. Functions that interact with the remainder of the VPP, that runs in the normal datatype, have to be wrapped to ensure seamless communication. Special entry points must be implemented that can call the lifting line method with standard datatype for normal VPP solving and the special datatype for gradient computation. For more information on ADOL-C see Walther and Griewank (2020).

The remainder of the VPP computes the V_{MG} influenced by the forces and moments acting on the yacht produced by the various components of the yacht (e.g. sails and foils). The total forces are the sum of the components forces. To link the design variables to the performance of the yacht, the effect of the forces and moments on the yacht´s performance must be included. This requires the computation of the gradient of V_{MG} with respect to $\boldsymbol{fm}$, so from 6 input variables to 1 output variable. Again, the number of inputs is higher than the number of outputs, which would generally favor the adjoint method, but the total number of inputs is small. Hence, the efficiency gain of the adjoint method is negligible while the implementation time can be reduced significantly by employing finite differences. The gradient is given by

$$\nabla V_{\mathrm{MG}}(fm) = \left[\frac{\partial V_{\mathrm{MG}}}{\partial f_x}(\boldsymbol{fm}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial f_y}(\boldsymbol{fm}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial f_z}(\boldsymbol{fm}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial m_x}(\boldsymbol{fm}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial m_y}(\boldsymbol{fm}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial m_z}(\boldsymbol{fm}) \right]^T \tag{11}$$

and is computed with a forward finite differencing scheme. Using the chain rule $[f(g(x))]' = f'(g(x)) * g'(x)$, the two parts can be linked together, where $g'(x)$ corresponds to $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ and $f'(g(x))$ to $\nabla {V_{\mathrm{MG}}}^T(\boldsymbol{fm})$. The resulting gradient is

$$\begin{aligned} &\nabla {V_{\mathrm{MG}}}^T(\boldsymbol{dv}) = \nabla {V_{\mathrm{MG}}}^T(\boldsymbol{fm}) * J_{fm}(\boldsymbol{dv}) \\ &= \left[\frac{\partial V_{\mathrm{MG}}}{\partial l_{ext}}(\boldsymbol{dv}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial l_{spa}}(\boldsymbol{dv}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial c_{ro}}(\boldsymbol{dv}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial r_{tap}}(\boldsymbol{dv}) \quad \frac{\partial V_{\mathrm{MG}}}{\partial \delta_{win}}(\boldsymbol{dv}) \right] \end{aligned} \tag{12}$$

and represents the influence of the hydrofoil design variables on the V_{MG} of the yacht. For example, the partial derivative of V_{MG} with respect to l_{ext} is hence given by

$$\begin{aligned} &\frac{\partial V_{\mathrm{MG}}}{\partial l_{ext}}(\boldsymbol{dv}) \\ &= \left(\frac{\partial f_x}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial f_x} + \frac{\partial f_y}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial f_y} + \frac{\partial f_z}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial f_z} + \frac{\partial m_x}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial m_x} + \frac{\partial m_y}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial m_y} + \frac{\partial m_z}{\partial l_{ext}} \frac{\partial V_{\mathrm{MG}}}{\partial m_z} \right). \end{aligned} \tag{13}$$

The gradient $\nabla V_{\mathrm{MG}}(\boldsymbol{dv})$ of the big parametric model is computed in the same manner but has 68 rows. $\nabla V_{\mathrm{MG}}(\boldsymbol{fm})$ remains similar. Splitting the gradient computa-

tion in two parts has a further benefit. The Newton-Raphson method solves for the steady sailing state in an iterative process. The valid steady sailing state is the last step of this process where all forces and moments equate to zero. $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ is only influenced by this final state of the boat, so the last step of the Newton-Raphson. A full differentiation using finite differences for example would however include one full solving process per input variable plus the initial. Hence, numerous states are computed that do not affect $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ and are not required for the approximation of $\nabla V_{\mathrm{MG}}(\boldsymbol{fm})$. With the split gradient computation this can be exploited. The steady sailing state is computed and then $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ is evaluated for the last step only. This requires little effort compared to a full solve, especially as the adjoint method is used. The gradient $\nabla V_{\mathrm{MG}}(\boldsymbol{fm})$ is then computed with 6 further VPP runs, which represents the influence of the forces and moments on the V_{MG}. This means 7 VPP runs plus the computation of $J_{\boldsymbol{fm}}(\boldsymbol{dv})$ and results in a significant further reduction in computational time. The Jacobians and the full gradients have been validated with finite difference approximations of gradients and showed excellent agreement. The computation of the gradient with the combined adjoint/finite difference method is thereby 36 times faster than a purely finite difference-based approach for 68 design variables.

8 General design optimisation

The first optimisation using the VPP-driven framework is conducted with the simple parametric model. The results of this optimisation are compared to the findings of the parametric study to ensure a correct implementation and set-up of the optimisation routine. In the parametric study, 72 different designs were tested comprising of any combination of the values listed in the "Values"-column in.

Table 3. The V_{MG} of these candidates were computed for up- and downwind conditions in 8, 11 and 14 kts V_{TW}.

For the comparison the upwind condition in 11 kts V_{TW} was selected, the results of which are shown in Fig. 10. In the upwind conditions in 11 kts V_{TW}, candidate 15 with the lowest root chord (0.4 m) and taper ratio (0.25) performed best. The small half wings result in low-wetted area foils as a larger portion of the weight is stored in the bulb. The round shape of the bulb has less wetted surface area per weight. This is accompanied by the lowest extension length (0.6 m), which again reduces wetted surface area and hence profile drag. The induced drag is reduced by maximising the span (1.88 m). The wing rake of 4 degrees is faster than the other settings. The influence is however negligible as the design variable wing rake is nearly identical to the control variable rake which is used for balancing the forces in z-direction. Candidate

Table 3 Comparison Set-up

Parameter	Values	Initial Design
Extension length l_{ext}	0.6 m, 0.95 m, 1.3 m	0.95 m
Semi span l_{spa}	1.75 m, 1.88 m	1.815 m
Root chord c_{ro}	0.4 m, 0.8 m	0.6 m
Taper ratio r_{tap}	0.25, 0.5	0.375
Wing rake δ_{win}	0.0°, 2.0°, 4.0 °	2.0°

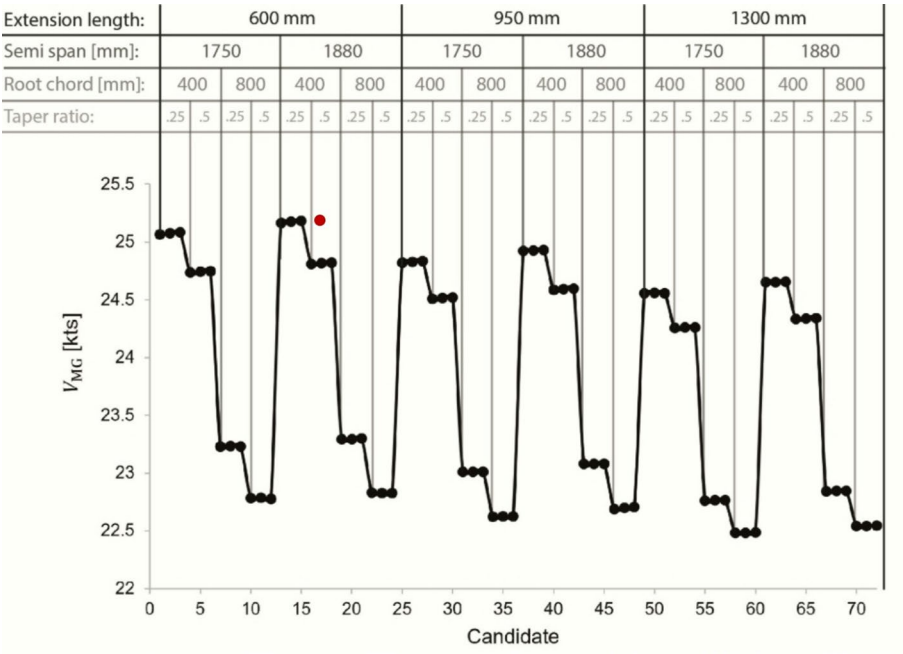

Fig. 10 Foil performance in 11 kts at $\beta_t = 47.5°$, **candidate 15 with highest** V_{MG} **in red**

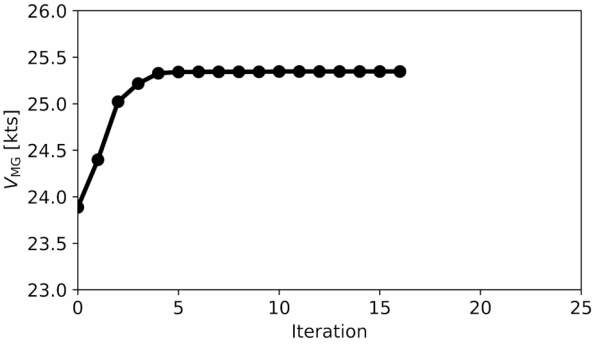

Fig. 11 Convergence history of the general design optimisation in 11 kts upwind

15 was 12% faster than candidate 70, the slowest in the study, and demonstrates the immense influence of hydrofoil design on yacht performance.

The corresponding adjoint optimisation run with the simple mode is started from the middle of the value ranges in Table 3 which is detailed in the "Initial Design"-column and shown in Fig. 12(a). The lowest and highest values of every variable are used as bounds. The convergence criterium is set to 1×10^{-4} in conjunction with an acceptable level criterium stopping the run if the changes are below 1×10^{-3} for five consecutive iterations. The maximum number of iterations is limited to 25.

The optimisation run converges to the optimal solution within 16 iterations as shown in Fig. 11. From a practical point of view however the optimal foil is already found at the fifth iteration after which the further design changes are negligible. The

final foil is presented in Fig. 12(b) and exactly corresponds to candidate 15, the fastest candidate from the parametric study. This proves the correct functioning of the approach. The optimum foil is 1.44 kts and hence 6.02% faster than the initial foil. This corresponds to a time saving of 12.8 s on a 1.5 nautical mile upwind leg. Two more optimisation runs were started, one from the lower bounds and one from the upper bounds. Both runs again identify the design corresponding to candidate 15 as optimal. The runs are solved to the acceptable level in 12 and 16 iterations, respectively. This suggests a uni-modal design space and reinforces the results from the initial run. A pure finite differences-based run found the same result. Due to the low number of input variables both methods were equally fast in this case and solved in a matter of minutes on a standard desktop pc.

The 36th America´s Cup was won by ETNZ with a yacht that was clearly the fastest in the fleet. The boat featured a T-foil with no anhedral and this foil is believed to have played a major role in the yacht´s superior performance. This is in contrast to the findings of the parametric study that suggests that the higher the anhedral angle is, the better the foil performs. The reason that the ETNZ´s T-foil performed better than the Y-foils is that it allowed the use of a single flap and a single control system. This meant that they could not change the direction of the force vector anymore, but it significantly reduced the required volume and wetted surface area. This possibility of reducing the required volume due to a single flap on the T-foil was not considered in the parametric study. The study is based on a constant density representative of a foil with two flaps/control systems. For a constant density foil with two flaps, the maximum anhedral foil is still superior as it has the lowest wetted surface area for the given density and is hence identified correctly as the optimal foil. A further benefit of the T-foil is that it can be used as a surface piercing foil at higher speeds which allows the reduction of wetted surface area when not needed. This is not possible for Y-foils that pierce the surface at a smaller angle which promotes ventilation. As ventilation cannot be modelled and the foil wing is assumed fully immersed at all times, this factor is also not influencing the optimisation towards a lower anhedral.

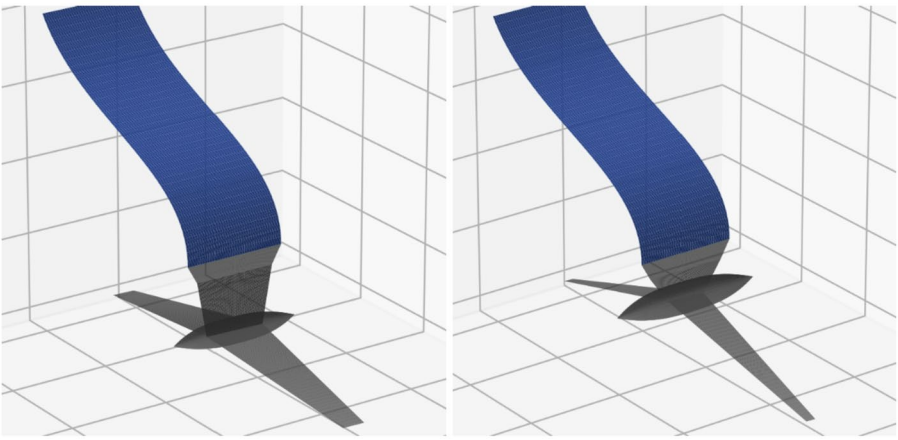

Fig. 12 Initial foil (**a**), optimal foil corresponding to Candidate 15 (**b**)

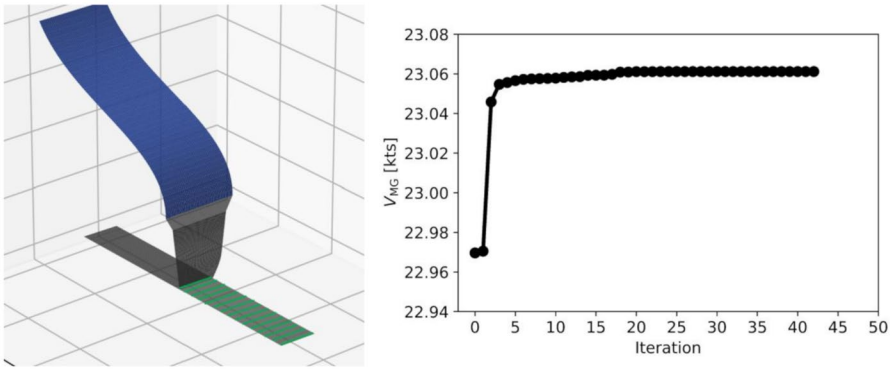

Fig. 13 Base foil for the twist-distribution optimisation, green lines represent the twist variables (**a**), convergence history of 11 kts upwind twist-distribution optimisation (**b**)

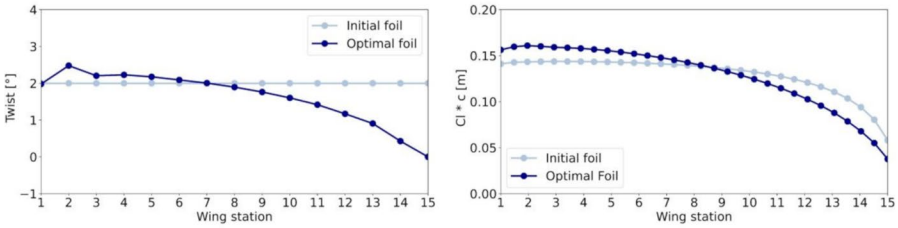

Fig. 14 Twist-distribution (**a**) and lift-distribution (**b**) of the initial and the optimal foil

9 Twist-distribution optimisation

According to Prandtl (1921), the induced drag of a hydrofoil is minimal when the lift-distribution is elliptic across the span of the foil. This can be achieved through a tailored twist-distribution. To validate whether the optimisation framework is capable of identifying the correct twist-distribution, an optimisation of the twist-distribution across 15 stations is performed with the aim to maximise V_{MG}. The optimisation is commenced for a straight T-foil without taper shown in Fig. 13 (a). The green lines represent the 15 sections where the twist can be modified by the optimisation algorithm on the outboard half wing. The inboard half wing is affected in the same way to fulfill the symmetry requirement of the foil wing. The convergence criterium is set to 1×10^{-4} in conjunction with an acceptable level criterium stopping the run if the changes are below 1×10^{-3} for five consecutive iterations. The maximum number of iterations is set to 50 accounting for the higher complexity of the problem compared to the general design optimisation.

The optimisation run converges after 42 iterations to the optimal solution as shown in Fig. 13 (b). The twist-distribution of the initial and the optimal foil are shown in Fig. 14 (a) and the resulting lift-distributions in Fig. 14 (b). The initial foil has a constant twist of 2° which, in conjunction with the constant chord and the absence of sweep, leads to the rectangular lift-distribution also observed in Figs. 7 and 8 (a).

This requires a reduction of the lift towards the tip and an increase of the lift towards the junction to achieve the desired elliptical lift-distribution. The optimal foil shows this characteristic and the resulting elliptical lift-distribution and proves that the optimisation framework is able to correctly optimise the twist-distribution for minimum induced drag. The optimal foil is 0.1 kts or 0.44% faster in terms of V_{MG}. Two more optimisation runs have been started, one from a twist-distribution with a decreasing twist towards the tip and one from a distribution with an increasing twist towards the tip. Both runs identified the same optimal twist-distribution as the initial run and solved to the acceptable level within 39 and 47 iterations, respectively. The runs solved in roughly 30 min on a standard desktop pc.

10 Detailed optimisation

Following the successful optimisations of the general design of the foil and the twist-distribution, a detailed optimisation is performed. The detailed optimisation uses the large parametric model with 68 design variables which allows the optimisation of the general design, the twist-distribution and a detailed spine- and plan form in a single run. The large, hybrid adjoint/finite-difference gradient was validated with the full finite difference gradient showing excellent agreement between the two methods. The computation of the gradient using the hybrid method is thereby 36 times faster than the full finite difference method. This demonstrates the immense efficiency of the developed approach.

The detailed optimisation is also started from the initial geometry shown in Fig. 12a. The same limits on extension length, semi span and chord along the wing apply (smallest allowed chord at every station corresponds to the chord at this location with minimum root chord and minimum taper ratio in the small parametric model). This allows validation with the previous results and a fair comparison between the general design optimisation and the detailed one. The model has additional freedom in the x- and z-position of every station (defined through the respective leading-edge point). The x-values can be manipulated by +- 0.2 m and the z-values can move between the bottom of the box and the waterline in the canted case. Bounding the foil at the waterline rather than the top of the box ensures that all designs are fully submerged. This still allows foils such as candidate 15 with the minimum extension length. Furthermore, the twist at every station is open to optimisation with +- 3°. The chord of the extension can be changed by +- 0.2 m from the initial design. The optimisation run is subject to the same convergence criteria as the previous runs with a maximum number of 150 iterations to account for the much higher dimensionality. The convergence history is shown in Fig. 15.

The optimisation run terminates with the maximum number of iterations of 150 reached. The routine fails to fully converge to the given criteria within this limit. A higher number of maximum iterations has also not lead to convergence, which suggests that the gradients are not perfectly accurate or the function has a very flat valley. The convergence history also shows small dents in the performance at for example iteration 120, which further suggests that the gradients are not perfectly accurate in some cases. These slight inaccuracies are introduced in the second part of the gradi-

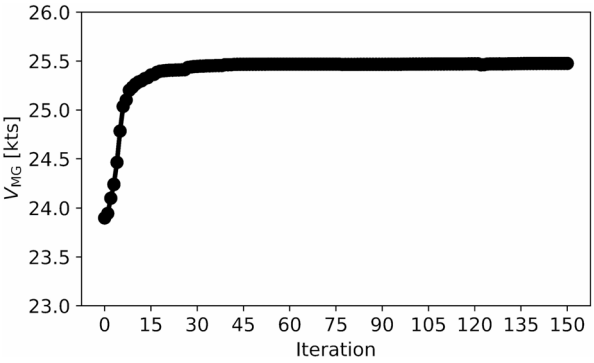

Fig. 15 Convergence history of 11 kts upwind optimisation, detailed model

ent computation. Here, the accuracy of the gradient is depending on the convergence of the Newton-Raphson determining the sailing state. Different step sizes and convergence criteria for the Newton-Raphson have been tested and the best setting has been used for the above run. Other settings have led to worse convergence in the optimisation. The step size of the gradient computation of $\nabla {V_{\mathrm{MG}}}^T(\boldsymbol{fm})$ also plays a role. A delta of 0.001 N and 0.001 Nm has delivered the best results. However, despite the lack of perfect convergence, the changes in the design towards the end of the optimisation are negligible from an engineering perspective. Hence, the convergence is deemed sufficient.

The final foil is 1.58 kts and hence 6.6% faster than the initial foil in terms of V_{MG}. For an upwind leg of 1.5 nautical miles, this results in 14 s time saving over the initial foil. Compared to the optimal foil of the general design optimisation, the detailed optimisation has increased the performance by a further 0.58% or 0.139 kts V_{MG}. Over a single race consisting of six legs, this results in a 7.2 s timesaving or a 94 m lead (4.1 boat lengths). This is a very significant advantage in this highly competitive event, where every fraction of a second and every meter counts. Having a faster yacht additionally has strategic and tactical advantages that are likely to amplify these performance gains.

The final foil has the same extension length and semi span as Candidate 15 which demonstrates the correct functioning of the routine with respect to the general design. The chord at all 15 half wing stations is minimised to the lower bounds. This results in the same root chord and taper ratio as Candidate 15 which was also at the lower bounds and proves that the method can find the correct results for high numbers of design variables.

The chord at the stations of the extension is reduced as much as possible to further reduce wetted surface area. The bulb is therefore slightly larger. This was expected and is reasonable. The final foil is shown in Fig. 16.

The view from the top in Fig. 16 shows that the final foil of the detailed optimisation is slightly swept aft and has a curved leading edge. This was not directly expected, but appears to help in achieving the elliptical lift-distribution shown in Fig. 17b. A separate optimisation run where the sweep was prohibited, however, also achieved an elliptical lift-distribution through a slightly higher twist towards the tip and a slightly lower one close to the junction. Therefore, it is not clear why the sweep should be beneficial. However, it is striking how smooth the swept and curved lead-

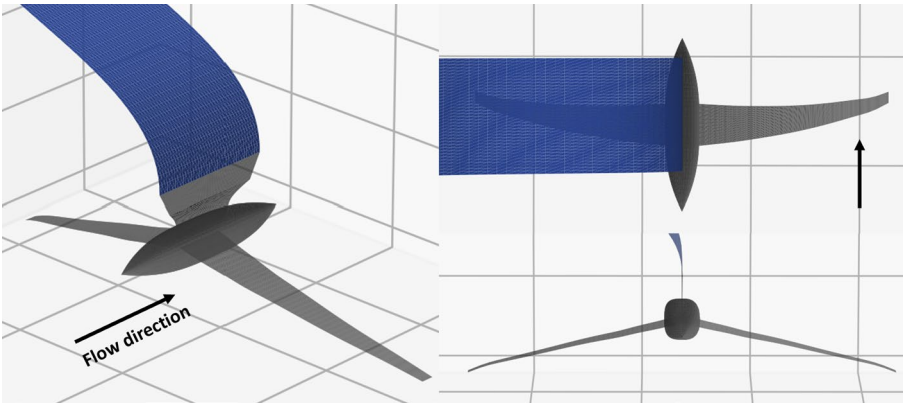

Fig. 16 Final foil of the detailed optimisation

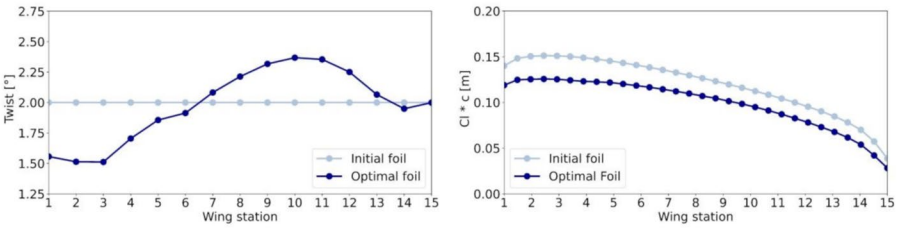

Fig. 17 Twist-distribution (**a**) and lift-distribution (**b**) of the initial and the detailed optimal foil

ing edge is. Whether this is only the optimum of the numerical model or also physically cannot yet be determined.

The twist-distributions of the initial and the final foils are shown in Fig. 17a. The initial foil has a constant twist of 2° which leads to a slightly triangular load distribution. The optimal foil has a lower twist towards the root and a higher twist towards the tip to counter this trend. The result is the desired elliptical lift-distribution that results in minimum induced drag. This is shown in Fig. 17b. The initial foil produces a higher lift as the optimised candidate when expressed in terms of $C_L * c$. This is due to the larger surface area of the initial foil. However, the optimised foil is resulting in higher speeds and the forces produced by the hydrofoil scale with the velocity squared. If the velocity is considered, both foils produce similar amounts of vertical lift and can support the yacht.

Several more optimisation runs have been started from different initial designs. This included different plan forms, spine forms and twist-distributions. In some cases, the same optimum as in the first run was found. In others, the foil designs were not able to achieve force equilibrium anymore after a certain number of iterations. In these cases, it is clear that the final foils are not the optimum. It shows that the results of such a detailed optimisation must be treated with caution. Nevertheless, the method has delivered consistent results from many other starting points reinforc-

ing the findings of the initial run. The optimisation runs solved in roughly 1.5 h on a standard desktop pc, demonstrating the efficiency of the method.

The results prove the capability of the method to conduct detailed design, while it is also capable of optimising the general dimensions such as the span and the extension length. This capability has resulted in a significant V_{MG} gain of 0.139 kts compared to the general design optimisation. The gains are expected to be even higher in optimisations for slower wind and hence boat speeds where the induced drag and the twist distribution has a higher influence. In addition, the detailed model allows further freedom to explore the design space, such as winglets for example, which was herein limited to enable validation and fair comparison with the general design optimisation. This shows the immense potential of the developed approach. Due to the high efficiency of the adjoint method, the framework can be extended to thousands of design variables which would allow to include the optimisation of hydrofoil sections and structural aspects.

11 Conclusions

Hydrofoils are the biggest performance differentiators in the "yacht" system and significantly affect the state and trim of the yacht. This requires the modelling of the entire system in an optimisation, which was achieved through the use of a stationary physics model of the entire yacht. It also requires detailed optimisation to fully extract the immense potential, which is prohibitively expensive with conventional optimisation strategies. This can only be achieved with a gradient-based optimisation algorithm paired with the adjoint method to compute the gradients. The adjoint version of the physics model was derived using the algorithmic differentiation tool ADOL-C for the bottleneck of the routine. The remainder of the physics model was differentiated with finite differences to reduce development time. The resulting routine computes the gradient of 68 input variables 36 times faster than the conventional full finite difference method and is the key to detailed system-based optimisation. To the authors knowledge, this makes the developed approach the first adjoint velocity prediction program for yachts and ships in general. The approach is not limited to hydrofoils on racing yachts but is beneficial whenever a component of a system has a huge influence on the system itself, provided that the model of the system is differentiable. The routine was first used to optimise the five variables of the parametric study, the results of which show excellent agreement to the findings of the optimisation. In a second step, a twist-distribution optimisation using 15 twist variables was conducted. The resulting twist-distribution led to the desired elliptical lift-distribution. Finally, a detailed optimisation of the entire foil with 68 design variables was performed. The results show that the method is capable of optimising the foil to great detail. The optimisation included the full spine- and planform as well as the twist distribution over 19 stations. The resulting foil has a minimum extension and minimum chord everywhere, paired with a maximum span and a twist-distribution that exhibits the elliptical lift-distribution. The detailed optimisation has increased the V_{MG} by 1.58 kts compared to the baseline foil. This is a very significant performance gain in the America´s Cup world. The optimisation run solved in 1.5 h on a standard desktop pc.

Due to this high computational efficiency, the framework can be extended to optimise thousands of design variables and would allow the inclusion of structural and section design. Further work is commenced to include additional true wind angles and true wind speeds to optimise the foil for a broader set of conditions.

Acknowledgements The authors are very grateful for the ideas and insights provided by Martin Fischer. They would also like to thank Studio Borlenghi and Stefano Gattini for allowing us to use their amazing photograph of the AC75.

Author contributions R.T. developed the adjoint system-based optimisation strategy and implemented it. He produced and analysed the results. He also wrote the main manuscript. The system-model is based on the software FS-Equilibrium, which was developed by K.H. The adjoint version of the bottleneck of the system-model was obtained through ADOL-C, an algorithmic differentiation tool partly developed by A.W. K.H. and A.W. also provided advise on how to use ADOL-C to differentiate FS-Equilibrium. S.R.T is the nominal supervisor and S.W.B the main supervisor to R.T´s. PhD project and provided general guidance. All authors reviewed the manuscript.

Funding This work was funded by the School of Engineering at the University of Southampton in the form of a PhD studentship and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy -- The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

Data availability Not applicable.

Declarations

Ethics approval and consent to participate Not applicable.

Consent for publication The authors consent to publish this paper in Springer's Optimization and Engineering journal.

Competing interests The authors declare no competing interests.

References

America´s Cup (2020) Virtual Eye, (viewed 8-Dec-2022). https://ac36.americascup.com/en/advanced-dashboard

Bingol O (2022) NURBS-Python Visualization, (viewed 5-Dec-2022). python.readthedocs.io/en/5.x/visualization.html

Drela M, Youngren H (2022) Athena Vortex Lattice, (viewed 10-Apr-2024). https://web.mit.edu/drela/Public/web/avl/

Eggert F (2018) Flight dynamics and stability of a hydrofoiling international moth with a dynamic velocity prediction program (DVPP). Technische Universität Berlin, Germany
Faltinsen O (2005) Hydrodynamics of high-speed marine vehicles. Cambridge University Press, Cambridge, UK
Garg N, Kenway G, Martins J, Young Y (2017) High-fidelity multipoint hydrostructural optimization of a 3-D hydrofoil. J Fluids Struct 71:15–39
Gattini S (2020) PRADA ACWS - Day 1 Photo Gallery, (viewed 22-Jun-2021). https://www.lunarossachallenge.com/en/gallery/614_PRADA-ACWS-Day-1-Photogallery
Griewank A, Walther A (2008) *Evaluating Derivatives - Principles and Techniques of Algorithmic Differentiation*. Society for Industrial and Applied Mathematics, 2nd Edition, Philadelphia, USA
Hansen H, Hochkirch K, Burns I, Ferguson S (2019) Maneuver simulation and optimization or AC50 class. J Sail Technol 4:142–160
Hart FP, Kriplani N, Luniya SR, Christoffersen CE, Steer MB (2006) Streamlined circuit device model development with fREEDAR® and ADOL-C. In: Bücker M, Corliss G, Naumann U, Hovland P, Norris B (eds) Automatic differentiation: Applications, Theory, and implementations. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 295–307
Hochkirch K (2018) *FS-Equilibrium user manual*
Hoerner SF (1965) Fluid-Dynamic drag. Published by the Author
Hospodář P, Drábek A, Prachař A (2022) Aerodynamic design and strength analysis of the wing for the purpose of assessing the influence of the bell-shaped lift distribution. Aerospace 9(1):13
Hovland P, Bischof C, Spiegelman D, Casella M (1997) Efficient derivative codes through automatic differentiation and interface contraction: an application in biostatistics. SIAM J Sci Comput 18(4):1056–1066
Jayaraman P (2022) GitHub – TinyNURBS, (viewed 31-May-2022). https://github.com/pradeeppyro/tinynurbs
Jee KW, McShan DL, Fraass BA (2005) Implementation of automatic differentiation tools for multicriteria IMRT optimization. In: Bücker M, Corliss G, Hovland P, Naumann U, Norris B (eds) Automatic differentiation: Applications, Theory, and implementations. Springer
Katz J, Plotkin A (1991) Low-speed aerodynamics: from wing theory to panel methods. McGraw-Hill, New York
Kim JG, Hunke EC, Lipscomb WH (2006) Sensitivity analysis and parameter tuning scheme for global sea-ice modeling. Ocean Model 14:1–2
Melis FM, Hansen H, Fischer M, Abdel-Maksoud M (2022) Velocity prediction program for a hydrofoiling lake racer. J Sail Technol 7:255–275
Meneghello G, Beyhaghi P, Bewley T (2016) Simulation-based optimization of the hydrofoil of a flying Catamaran. Massachusetts Institute of Technology, Cambridge, US
Molland AF, Turnock SR (2022) Marine rudders, hydrofoils and control surfaces: principles, data, design and applications. Butterworth-Heinemann an imprint of Elsevier, Oxford, Cambridge
Offshore Racing Council (2001) International Measurement System
Paulin A, Hansen H, Hochkirch K, Fischer M (2015) Performance assessment and optimization of a C-class catamaran hydrofoil configuration, 5th High Performance Yacht Design Conference Auckland, New Zealand
Phillips WF, Snyder OD (2000) Modern adaptation of prandtl's classic lifting-line theory. J Aircr 3:662–670
Piegl L, Tiller W (1997) The NURBS Book, 2nd edn. Springer Berlin / Heidelberg.
Ploe P (2018) Surrogate-based optimization of hydrofoil shapes using RANS simulations. École centrale de Nantes, France
Prandtl L (1918) *Tragflügeltheorie.* Königliche Gesellschaft der Wissenschaften zu Göttingen, Göttingen, Germany
Prandtl L (1921) *Applications of modern hydrodynamics to aerodynamics: NACA Report No. 116*
Röbenack K (2007) Controller design for nonlinear multi-input – multi-output systems based on an algorithmic plant description. Mathematical and Computer Modelling of Dynamical Systems. 13(2):193–209
SumToZero (2021) *GOMBOC Designer*, (viewed 31 January 2021). http://sumtozero.com/products/gomboc-designer/
Tadjouddine M, Forth SA, Keane AJ (2006) Adjoint differentiation of a structural dynamics solver. In: Bücker M, Corliss G, Naumann U, Hovland P, Norris B (eds) Automatic differentiation: Applications, Theory, and implementations. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 309–319

Tannenberg R, Turnock SR, Hochkirch K, Boyd SW (2023) VPP driven parametric design of AC75 hydrofoils. J Sail Technol 8:161–181

Wächter A, Biegler L (2006) On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. Math Program 106(1):25–57

Walther A, Griewank A (2020) *ADOL-C: a package for the automatic differentiation of algorithms written in C/C++*.

Authors and Affiliations

Rafael Tannenberg[1] · Karsten Hochkirch[2] · Andrea Walther[3] · Stephen R. Turnock[1] · Stephen W. Boyd[1]

✉ Rafael Tannenberg
rafael.tannenberg@soton.ac.uk

[1] University of Southampton, Southampton, UK

[2] DNV Ship Performance Center, Hamburg, Germany

[3] Humboldt-Universität zu Berlin, Berlin, Germany