XI International Conference on Computational Methods in Marine Engineering

Optimizing Wing Sails with Actuator Line Simulations and an Effective Angle of Attack Controller

Jarle V. Kramer[1,*]

[1] Norwegian University of Science and Technology,
Department of Marine Technology, Professor J.H.L Vogts veg 1a, 7052 Trondheim, Norway.

[*] jarle.a.kramer@ntnu.no

ABSTRACT

This paper presents two ways to simplify the optimization of the angle of attack of multiple wing sails: the actuator line method and an effective angle-of-attack controller. The actuator line method reduces the computational time compared to full CFD, while capturing interaction effects between sails and the rest of the ship. The effective angle of attack controller adjusts sail orientation based on local flow measurements, as an alternative to a full optimization procedure. The rationale for this controller is based on the observation that the optimal angle of attack is usually the one that maximizes the lift. This combined approach enables practical optimization of multi-sail wind propulsion systems. Case study results on a 200-meter bulk carrier demonstrate significant sail-ship interaction effects, with optimal orientations differing substantially from those predicted by isolated sail models.

Keywords: wing sail; aerodynamic interaction; actuator line method; effective angle of attack; lifting-line method; optimization.

1 INTRODUCTION

To achieve significant fuel savings on merchant ships using modern sails, multiple units are typically installed as several sails allow for a larger total sail area. However, optimizing multiple sails is more challenging than optimizing a single sail. This paper focuses on wing sails, where the operational variables are the angle of attack and flap angle. For maximum thrust applications, the flap angle can be set to maximize lift, making the angle of attack the primary variable for optimization. The challenge is that local flow around each wing sail varies significantly due to lift-induced velocities and interaction effects with the deck, equipment, and superstructure. The operation of all sails must therefore be evaluated considering these flow disturbances. Setting the angle of attack based on theoretical apparent wind conditions and single-sail optimization may not yield optimal results in practice.

Simulating multiple sails with the complete ship geometry is significantly more time-consuming than simulating a single sail. Additionally, multi-variable optimization requires several objective function evaluations, making full Computational Fluid Dynamics (CFD) simulations impractical for optimization loops. One solution to this is to use simplified and computationally faster simulation methods. One common approach is to use methods based on potential theory, such as the discrete lifting line method. Examples of such approaches can be found in Kramer and Steen (2022) and Malmek et al. (2024). These methods have several benefits; they are computationally efficient, robust, and easy to implement. Sail-to-sail interaction effects can be captured, and viscous effects on both lift and drag are easily accounted for. However, they cannot model the interactions with the rest of the ship, at least not directly. It is possible to use velocity fields from external sources, such as CFD simulations,

as input to a lifting line method, thereby capturing some interaction effects. However, this approach does not capture situations where the sails themselves affect the flow over the deck and superstructure.

This paper examines two methods designed to facilitate the rapid optimization of the angle of attack for wing sails, particularly when interaction effects are significant. The first method is the actuator line method. The primary benefit of this method is its ability to capture the interaction between the sails and the rest of the ship in a direct manner. That is, the flow around the superstructure affects the forces on the sails, and the lift-induced velocities from the sails also affect the flow around the superstructure. Initially developed for wind turbine interactions (Sørensen and Shen, 2002), it has been applied to various lifting surfaces, including sailboat keels (Renzsch and Ward, 2021) and wing sails (Kramer and Steen, 2023). The second method is a simplified controller that adjusts the sail orientation based on local effective angle-of-attack measurements along the wing span. This alternative to full optimization requires significantly less computational time and can still provide operational values that maximize thrust for most wind directions.

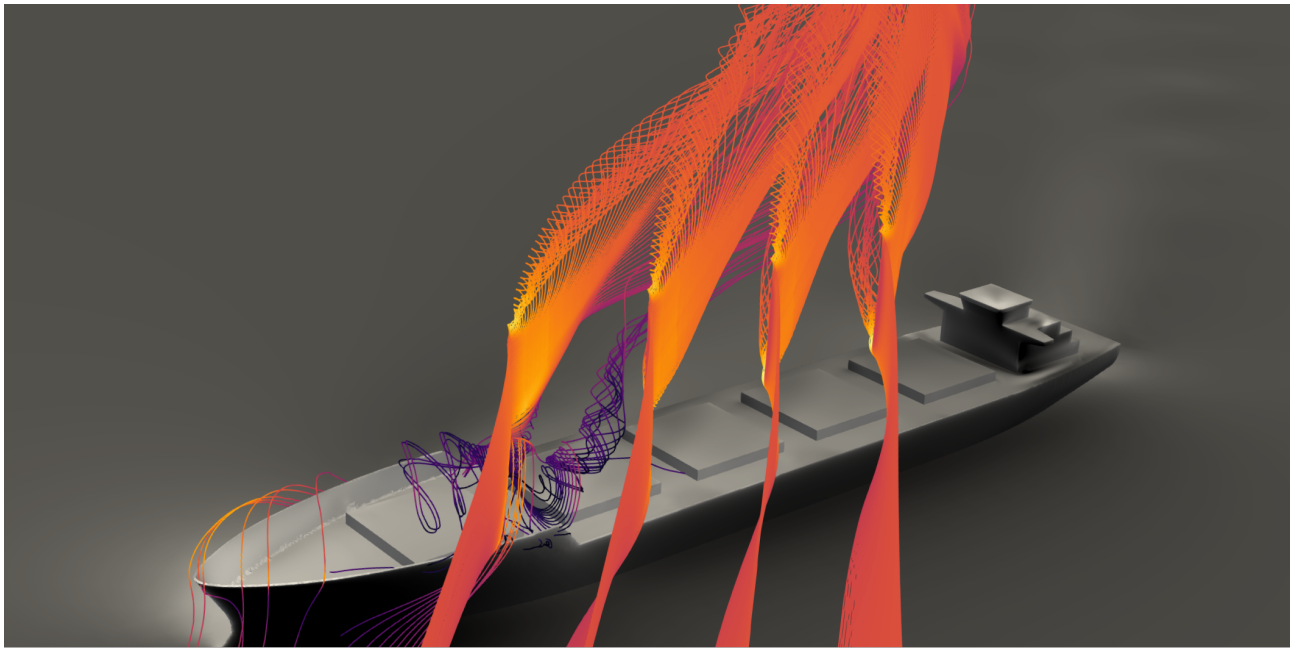

Figure 1: Visualizations of streamlines from line sources placed the locations of the sails from an actuator line simulation

2 CASE STUDY

The focus is on the aerodynamics and tuning the sails to produce the maximum thrust. This simplifies several other factors that may influence the optimal operation of sails, such as the negative hydrodynamic effects due to side forces, yaw moment, and heel moment, as well as the effect of variable thrust on both the propulsion machinery and propeller. See, for instance, Kramer and Steen (2022), for more on the importance of hydrodynamics effects on the sail control. Maximizing the thrust is a good starting point for further optimization later.

The ship in this study is the 200-meter-long SINTEF Ocean Bulk Carrier. More information can be found in SINTEF Ocean (2024). The hull above water was taken directly from the official geometry published online; however, since the ship does not have an official public superstructure geometry as of the time of writing, a simplified geometry was explicitly created for this case study. An illustration of the ship is shown in Fig. 1. The sails consist of four two-element wing sails, with the same design as that used in the recent comparison study of route simulations, known as SWOPP part 2 (no public

source available). The chord length of the main element is 4 m, the chord length of the flap is 2 m, and the span is 23 m. The flap angle is set to 15 degrees. The sails are placed along the center line of the hull, at longitudinal positions equal to 38.0 m and 10.5 m behind the midship and 17.5 m and 46 m ahead of the midship. The sail geometry is visualized in Fig. 2.

Throughout the paper, force results are presented as non-dimensional coefficients. These are always calculated by dividing any force value by the dynamic pressure based on the inlet conditions and the relevant sail area. Lift and drag refer to forces normal and parallel to the inlet velocity, while thrust and side force refer to forces pushing along and normal to the ship's center line. These values are generally presented as a function of the wind direction. The atmospheric boundary layer is not modeled, for simplicity, so the wind direction corresponds to an ideal uniform apparent wind direction.

3 METHODS

Three simulation methods are employed: a discrete lifting line method, the actuator line method, and conventional CFD simulations. The lifting line and actuator line methods share much of their underlying functionality. As such, the implementation of the methods used in this paper is based on a library designed to unify as much as possible between the different methods.

3.1 Lifting line

The lifting line method represents the wings as several discrete vortex line elements both along the span of the wing and the wake. That is, the wings are represented as several discrete horseshoe vortices where the trailing vortices follow the freestream and the bound vortices are located at the wing span line. For each line element making up the span of the wings, there is also a control point and a chord vector. The induced velocities from the vortex elements in the simulations are evaluated at the control point, and the angle of attack is calculated relative to the chord vector. The circulation strength on each line element is then calculated from a parametric two-dimensional *sectional force model*, the angle of attack, the velocity, and the chord length. The sectional force model is the only empirical part of the model, and it needs to be tuned to the specific sail. The main parameters to be set are the drag and lift at zero angle of attack, assuming no lift-induced velocities (i.e., the 2D lift and drag), and the angles of attack at which the foil will stall. The output from the sectional model, as used in this case study, is shown in Fig. 3. As there is a coupled relationship between the circulation strength and the induced velocities from the wake, it is necessary to use a solver to find the right value. This is achieved with a simple and robust solver, where the strength is gradually updated with the most recent estimations of the velocity field, utilizing a damping factor in the update to stabilize the solution.

3.2 Actuator line

The actuator line method uses the same form of line elements and sectional model as the lifting line method to represent the wings themselves. However, an actuator line does not have a potential theory wake. Rather, the model is made to be used within a CFD solver, and the velocity at each control point is evaluated by interpolating the velocity directly from the velocity field in the simulation. Lift-induced velocities from the sails are modeled by projecting the calculated forces from the line elements back into the CFD solver using volume forces applied to the momentum equations. How the forces are projected into the CFD domain can significantly affect the strength of the lift-induced velocities. This case uses an anisotropic Gaussian projection, inspired by the work in Churchfield et al. (2017), where the length parameter in the Gaussian distribution is set to 0.4 times the chord length in the

chord direction and 0.15 times the chord length in the thickness direction.

3.3 CFD setup

The CFD simulations, including the actuator line simulations, were performed with OpenFOAM (OpenCFD Ltd, 2025), using an unsteady incompressible solver and the k-omega SST turbulence model. The mesh design is partially visualized in Fig. 2. The resolution around the wing sails was set to 256 cells per chord length and six inflation layers, which corresponds to a theoretical y+ value of 45. The actuator line simulations were executed with 32 cells per chord length in refinement zones around the line geometry itself. This resolution was based on a convergence study performed in Kramer and Steen (2023). The mesh resolution around the deck and superstructure was set to vary between 100 and 200 cells per ship length.

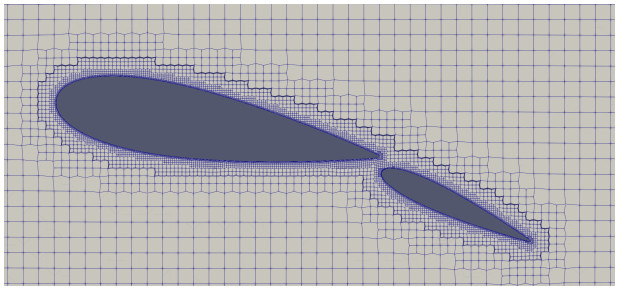

(a) Illustration of the mesh around a section of the wing sail

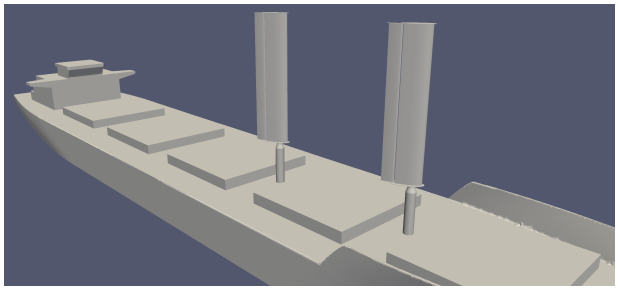

(b) Illustration of wall patches making up the wing sail and ship for the validation case

Figure 2: Visualization of the CFD geometry

3.4 Comparison of methods

The lifting line method has previously been compared against CFD in Kramer and Steen (2022), and both the lifting line method and actuator line method have been tested in Kramer and Steen (2023). In addition, two new tests are done for this paper to evaluate the differences between the simplified models and a full CFD simulation. The first test checks the model's ability to compute forces on a single sail, shown in Fig. 3. This serves two purposes: first, the result is used when tuning the two-dimensional sectional model used as input to both the lifting line and actuator line simulations. This is achieved by varying the model's parameters to match the output from the lifting line simulations against the CFD results. Second, it is used to check how accurately the models can estimate the lift-induced velocity from a single sail. Note that the actuator line predicts slightly lower lift-induced drag than both the CFD and the lifting line method, but achieves a good match in lift. This is fairly typical for actuator line simulations, and a possible solution is to add a small empirical correction to the lift-induced drag. However, the result presented in this paper is shown without such a correction.

The second comparison case involves simulating both the ship and sails together, using both the actuator line approach and full CFD. To reduce the computational load for this comparison study, only the two front sails were used, assuming the physics for two sails to be representative also for four sails. The result is shown in 4, showing a relatively good match between the methods. Note that the figure illustrates relatively significant interaction effects between the sails and the ship, and that the trend in forces from both methods is the same. For one, the front sail, labeled *Sail 2* in the plot, is experiencing significantly less lift than the stern sail. The opposite would be the case if only sail-to-sail interaction effects were present. Second, the lift is reduced for both sails when the apparent wind direction is increased. At least for the stern sail, this would not be the case if the sails were standing alone.

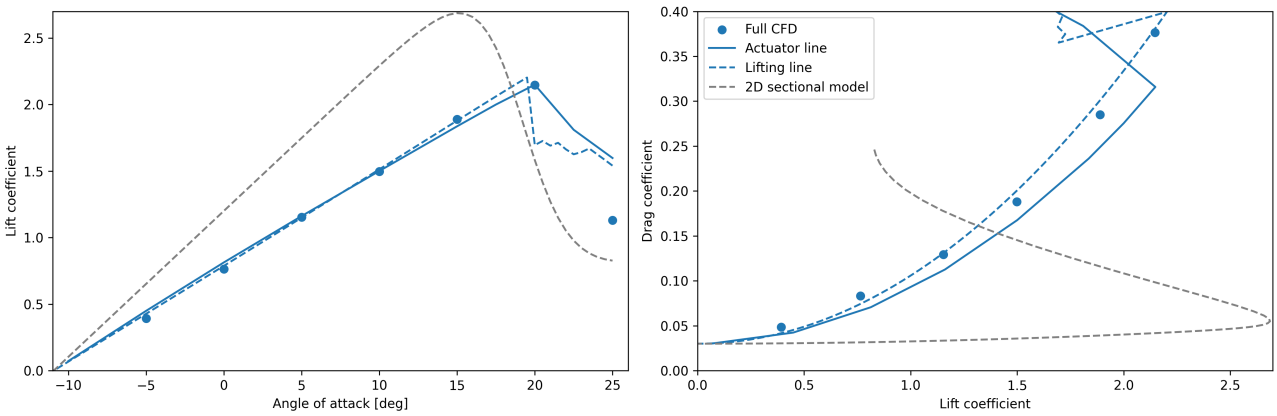

Figure 3: Comparison of lift and drag from lifting line, actuator line, and full CFD simulations. The sectional model used as input to the lifting line and actuator line method is also shown.

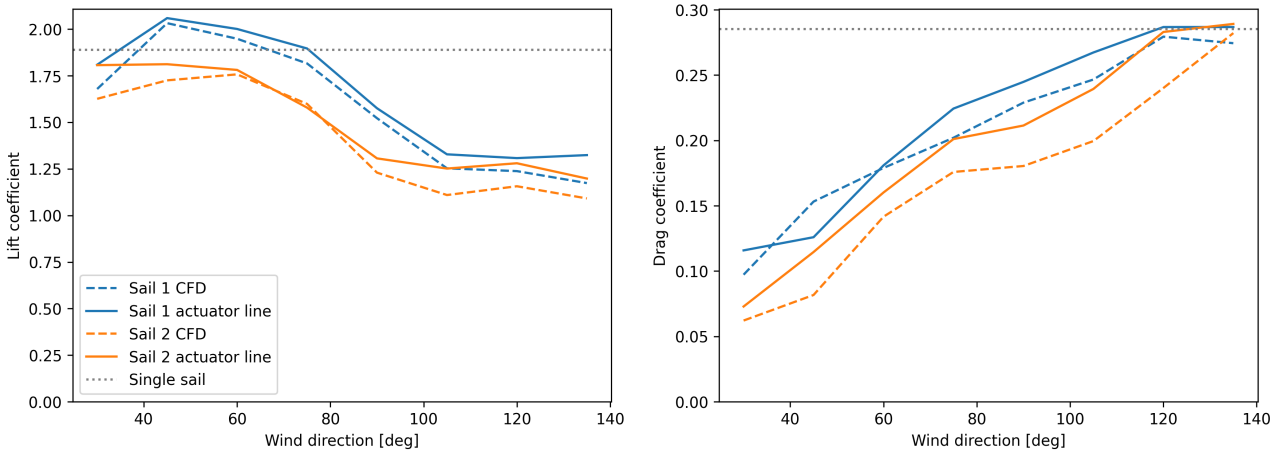

Figure 4: Comparison between the actuator line method and full CFD simulations for two sails at the front of the ship, placed with an angle of attack of 10 degrees relative to the free stream.

3.5 Effective angle of attack controller

Although actuator line simulations are significantly faster than conventional CFD, they remain relatively slow for cases requiring optimization for multiple wind directions, speeds, and ship velocities. The exact computational time depends on the computer used; however, some rough values for a 56-core computer are provided below. A lifting line simulation takes tens of milliseconds, while actuator line simulations require around 10 minutes for a single sail. Simulating several sails on a ship can take hours, depending on the resolution used for the ship geometry. Although a comparative conventional CFD simulation would take days, making the speed-up of the actuator line method significant, it would still be beneficial to avoid running a formal optimization procedure for every case.

The effective angle of attack controller adjusts the sail orientation so that each sail achieves the same maximum local angle of attack value during a single simulation, thereby reducing the need for multiple iterations with an optimization algorithm. The procedure is first to calculate the difference between the measured maximum angle of attack and the controller set point, and then adjust the sails accordingly with a 0.5 damping factor to ensure stability. Updates occur every 40 time steps to minimize computational overhead associated with recapturing relevant cells for force projection. This approach was inspired by lifting line optimization results that account for sail-to-sail interaction. Previous work on this topic is presented in Kramer and Steen (2022), and some additional results

for this case study are illustrated in Fig. 5. When maximizing thrust, the optimal angle of attack is typically the angle that maximizes lift. The notable exception is close-hauled conditions. The maximum lift for a wing sail occurs at the stall angle. However, operating the sails exactly at this angle might be risky for real-life operation. For one, the exact stall angle is highly uncertain, as it is difficult to predict accurately with simulations. Second, the stall angle might depend on both the Reynolds number and the turbulence level in the atmosphere, and will therefore vary over time. Third, wind gusts and the ship's motions will cause the angle of attack to vary on the sails, both above and below the steady-state condition. As such, it seems reasonable to limit the effective angle of attack to a value below the stall angle.

A line search method is used to optimize the angles of attack with the lifting line method. A quadratic penalty in the objective function is used to enforce a maximum limit on the effective angle of attack. That is, the objective function is set to maximize the thrust, plus a quadratic penalty whenever the maximum effective angle of attack is above the limit. As seen in Fig. 5, for most wind directions, the result of such an optimization is that sails are generally set to operate at the orientation that gives the maximum effective angle of attack. The theory behind the effective angle of attack controller is that the same principle would apply even when the sails experience interaction effects with the rest of the ship.

4 RESULTS

Figure 5 shows the measured effective angle of attack along the span of the sails for two wind directions to show that the effective angle of attack controller performs its job of tuning the maximum angle of attack. In the case of the wind direction of 45 degrees, it is seen that the sails with the fixed orientation experience a larger effective angle of attack along the span, likely due to interaction with the deck and superstructure. The result of the controller in this case is a reduced angle of attack, which also results in a reduced lift compared to the fixed case. However, it then ensures the required margin against stall. In the case of 90-degree wind direction, it is seen that the effective angle of attack is lower with a fixed orientation than with the controller. In this case, the controller increases the angles, which gives higher values for the thrust. Note also that the sail furthest towards the bow experiences very different values than the other sails, likely due to interaction with the forecastle.

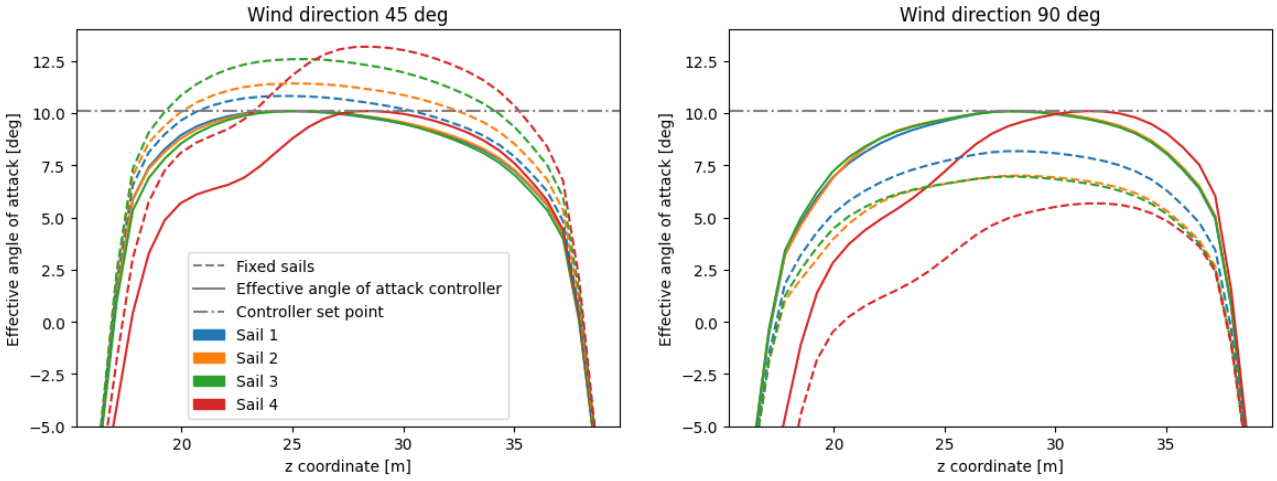

Figure 5: Comparison of the effective angle of attack along the sails with both fixed orientation of the sails and with an effective angle of attack controller where the set point is 10.1 degrees. The sail index starts at 1 at the sail closest to the stern and ends at 4 for the sail closest to the bow.

The thrust and side force from different simulations are shown in Fig. 6. Two set points for the effective

angle of attack controller are tested to determine their impact on the result. The lowest value was chosen based on the maximum effective angle of attack when a single sail is simulated with a 15-degree angle of attack. That is, the effective angle set point represents the same margin against stall as a single sail operated at 15 degrees angle of attack. For wind direction from 75 degrees and above, the controller can increase the total thrust from the sails. However, for smaller wind directions, there is little to no improvement. This is, as mentioned above, due to the larger angle of attack experienced by the fixed sails due to interaction effects. To demonstrate that it is possible to increase the thrust with the controller, it was also executed with 12.5 degrees as a set point, which roughly corresponds to the largest angle for the fixed sails in Fig. 7.

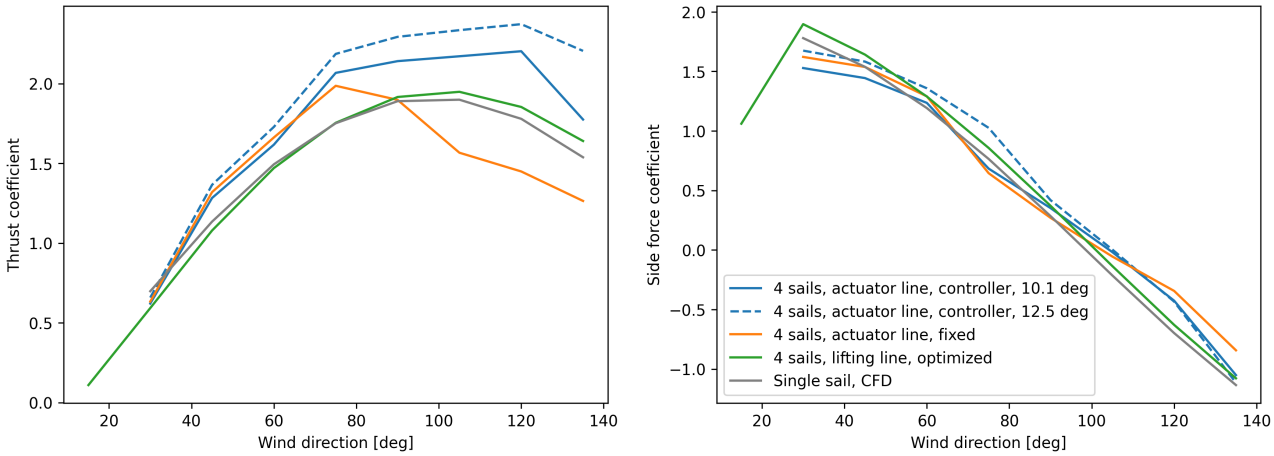

Figure 6: Thrust and side force with different control strategies for the sails

Finally, Fig. 5 compares the resulting angles of attack for the four sails from both lifting line simulations and from actuator line simulations with an effective angle of attack controller. The right plot also shows the measured effective angle of attack in the lifting line simulations when the geometric angle is optimized with a limit on the maximum value for the effective angle of attack, as explained in Section 3.5. This demonstrates that, when only sail-to-sail interaction effects are accounted for, the optimal geometric angle is primarily the angle that maximizes the effective angle, which is also the angle that maximizes lift. The left side of Fig. 5 shows that the resulting angles from the effective angle of attack controller are widely different from the optimal angle from the lifting line simulations. This is, again, an indication of quite strong interaction effects with the rest of the ship.

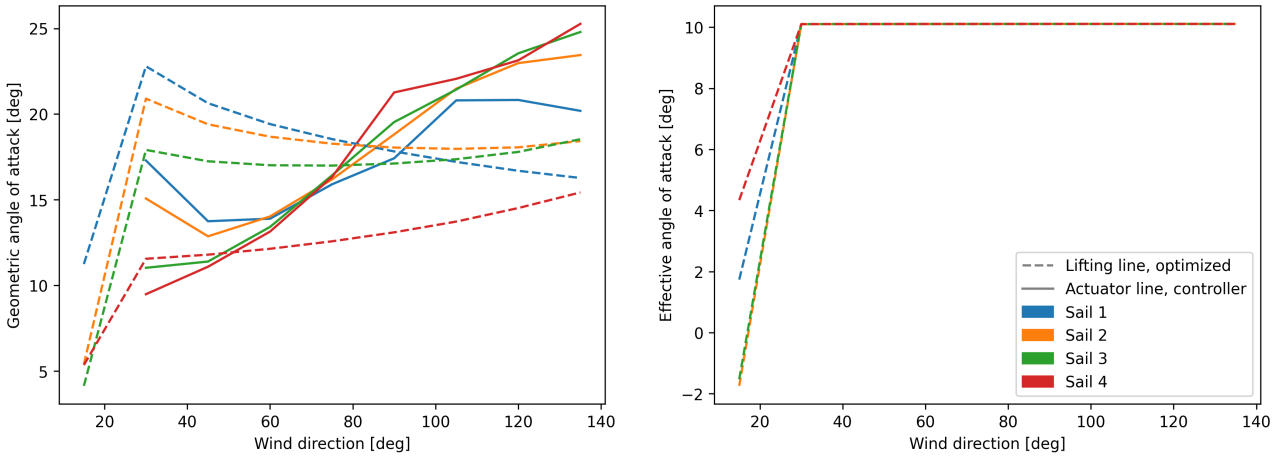

Figure 7: Comparison of the resulting angles of attack from the different control strategies. The sail index starts at 1 at the sail closest to the stern and ends at 4 for the sail closest to the bow.

5 CONCLUSIONS

The actuator line method is found to be efficient at simulating interaction effects between sails and the rest of the ship. However, the simulation time is still sufficiently long that it may become impractical to use conventional optimization algorithms directly. As such, the effective angle of attack controller was suggested as a simplified approach. The controller maximizes the lift by maximizing the effective angle of attack, and also enables control of the margin against stall. The simulations indicated relatively strong interaction effects between the sails and the rest of the ship. This was evident in both the forces, which deviated significantly from what was expected based on the simulated forces on a single sail in ideal conditions, and in the resulting angles from the controller. The results indicate that maximum lift and thrust are achieved at angles that differ significantly from those predicted by an optimization algorithm using a discrete lifting line model. The latter model includes sail-to-sail interactions, but neglects the ship-to-sail interaction. It therefore seems important that both the control policy and the force prediction are tested and tuned based on models that account for the full interaction effects. The combination of actuator line simulations and an effective angle of attack controller enables the control policy to be tuned during a single simulation, making it a practical approach. Further optimization may be necessary in some cases, such as accounting for hydrodynamic effects and performing additional fine-tuning when sailing in closed-haul conditions. However, the simplified controller can still be of use, both to compute the initial starting point for further tuning and, perhaps, to simplify optimization by adjusting the set-points of the controller. The latter is planned for future work on this topic.

ACKNOWLEDGMENTS

The author gratefully acknowledges the financial support from the Research Council of Norway granted through NFR project 344295, KSP-S WIND.

References

M. J. Churchfield, S. J. Schreck, L. A. Martinez, C. Meneveau, and P. R. Spalart. An advanced actuator line method for wind energy applications and beyond. In *35th Wind Energy Symposium*, page 1998, 2017.

J. Kramer and S. Steen. Sail-induced resistance on a wind-powered cargo ship. *Ocean Engineering*, 261, 2022. doi: 10.1016/j.oceaneng.2022.111688.

J. Kramer and S. Steen. Actuator Line for Wind Propulsion Modelling. In *Numerical Towing Tank Symposium*, Ericeira, Portugal, 2023.

K. Malmek, L. Larsson, S. Werner, J. W. Ringsberg, R. Bensow, and C. Finnsgård. Rapid aerodynamic method for predicting the performance of interacting wing sails. *Ocean Engineering*, 293, 2024. doi: 10.1016/j.oceaneng.2023.116596.

OpenCFD Ltd. OpenFOAM home page, 2025. URL https://www.openfoam.com/.

H. Renzsch and B. Ward. A RANS-BEM method to efficiently include appendage effects in RANS-based hull shape evaluation. *Journal of Sailing Technology*, 6(01):44–57, 03 2021. doi: 10.5957/jst/2021.6.3.1.

SINTEF Ocean. SOBC-1 web page, 2024. URL https://www.sintef.no/en/sintef-research-areas/maritimetransport/sintef-ocean-bulk-carrier/.

J. N. Sørensen and W. Z. Shen. Numerical modeling of wind turbine wakes. *Journal of Fluids Engineering*, 124(2):393–399, 05 2002. doi: 10.1115/1.1471361.