Parametric Optimization of a Kiteboarding Hydrofoil Using Computational Fluid Dynamics

AUTHOR

Zachary P. Backas
238 S. Kenmore Ave., Elmhurst IL 60126
(630) 464-9224 – zbackas@gmail.com

ADVISOR

Professor Adrian Onas
Webb Institute
298 Crescent Beach Rd.
Glen Cove, NY, 11542
Webb.edu

Zachary P. Backas
Webb Institute

Parametric Optimization of a Kiteboarding Hydrofoil Using Computational Fluid Dynamics

ABSTRACT

Kitefoils are a form of hydrofoil that enables a kiteboarder to ride above the surface of the water on the lift of the appendage alone. Kitefoils are designed to produce a lift and moment to support and stabilize a rider across a wide range of operating speeds. The interconnectivity between the aspects of geometry that are significant to a kitefoil's performance means that choosing the design parameters that maximize the foil's efficiency is a complex challenge. As such, this hydrodynamic problem lends itself well to a parametric optimization procedure, in which the performance-critical aspects of the hydrofoil geometry are varied in search of an optima. In order to analyze design variants in the optimization process, a computational fluid dynamics analysis is employed and coupled to the process for automated geometry variant analysis. Ultimately, the performance of a base hydrofoil geometry may be improved through parametric optimization, though the effectiveness of the optimization is dependent upon sufficient parameter refinement and design space exploration.

NOMENCLATURE

CFD	Computational Fluid Dynamics
AoA	Angle of Attack
RANS	Reynolds Averaged Navier Stokes
CAESES	Computer aided engineering system empowering simulation
A_m	Main foil Area
A_s	Stabilizer foil area
$C_{fuse\_extension}$	fuselage extension scale factor
C_{main_area}	Main foil scale factor
C_{stab_area}	Stabilizer area scale factor
$G_{fuse\_extension}$	fuselage extension gain factor
G_{stab_area}	Stabilizer area gain factor
L	Fuselage length
p_d	Dynamic Pressure
α	main foil root twist angle
ΔA	Stabilizer foil area change
ΔL	Fuselage length change
$\Delta\theta$	Main foil angle of attack change
θ	Main foil angle of attack

INTRODUCTION

Kitefoils are a type of hydrofoil designed to lift a kiteboarder out of the water. They are designed to decrease resistance and enable the kiteboarder to ride in lower wind, at higher speeds, and at more favorable angles to windward. New foil designs often challenge previous standards, prompting a need for a thorough analysis of kitefoil designs with respect to their varying forms. Numerous successful designs exist on the market today that share the same basic features but differ significantly in their performance characteristics. The complex geometry and design constraints associated with kitefoils lend themselves to optimization by way of a parametrized base-model. A parametrized base-model may be programmed with geometric variability that allows for the hydrofoil characteristics that influence performance to be easily varied within the optimization process.

OBJECTIVES

The objectives of this thesis are (1) to develop a working procedure for parametric optimization in the programs CAESES and Star-CCM+ and (2) to utilize that to improve the performance of a provided base-model kiteboarding hydrofoil through a parametric optimization process.

BACKGROUND

THE PERSONAL HYDROFOIL

Personal hydrofoils originate from the early 1960s when Walter Woodward, an aeronautical

engineer, developed the first waterski hydrofoil (Hydrofoiling History). The foil, marketed as the Dynaflite Hydrofoil, consisted of two water skis with a strut attached to each ski and bi-wing style lifting surfaces as shown in Figure 1.

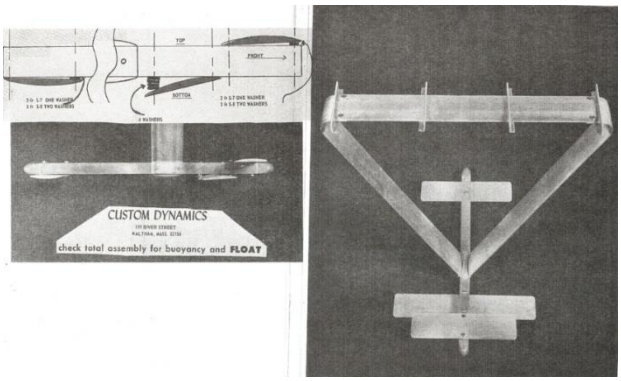

Figure 1. Dynaflite Hydrofoil
http://www.hydrofoil.org/history.html

The hydrofoil was designed as a bi-plane to facilitate a low take-off speed of around 10 mph. The design also included a stabilizing foil aft of the struts. The foil was adjustable in the sense that the larger of the two bi-plane wings could be removed to facilitate higher top speeds at the expense of a higher take-off speed.

KITEBOARDING HYDROFOILS

Based on the towed hydrofoils initially developed, the personal hydrofoil was applied to kiteboarding. Modern kiteboarding hydrofoil designs vary in their geometry based upon their intended use, typically falling into either beginner, freestyle, or race categories. Beginner kitefoils characteristically have relatively low-aspect foils that are designed to facilitate stable foiling at low take-off speeds. Freestyle foils are designed for jumping and maneuverability and are generally built to achieve moderate to high speeds with more structure than a race foil to accommodate landing impacts. Racing kitefoils are designed for course racing, which requires a balance of both upwind and downwind performance. The goal of a racing kitefoil is to allow a rider to round a course in the fastest possible time. Ultimately, the foil efficiency, ease of ride, and balance of upwind and downwind speed contributes to a successful race foil design. The high-performance nature of racing kitefoils results in designs that characteristically have high-aspect-ratio foils, minimal structure, and carbon fiber construction. A typical racing kitefoil with its main components defined is shown in Figure 2.

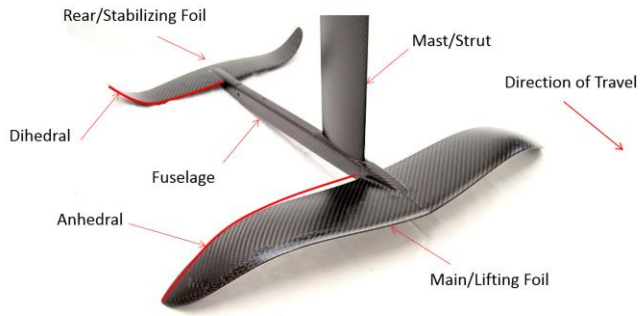

Figure 2. Kitefoil geometric features
Adapted from Spotz

OPTIMIZATION RATIONALE

Kitefoils are a combination of a number of individual components, each with a specific purpose and particular constraints that it is subject to. The juxtaposition of each part of a kitefoil into the ultimate hydrofoil system creates a complex geometry with each component's performance both affected by and an influence of the performance of the kitefoil as a whole. While it is easy enough to create a single lifting surface with given operating constraints based upon well-established airfoil design practices, the development of a system of hydrofoils for a complex range of operating conditions is a more complex problem that does not likely have a single solution. The complex tradeoffs between various kitefoil geometric features lend themselves well to an automated optimization process.

THEORY

KITEFOIL DESIGN

Kitefoils must produce enough lift to support a rider underway while also producing a moment of a manageable magnitude for the rider to balance off of. The kitefoil must be able to produce a sufficient amount of lift in a speed

range that extends from "take-off" speed to top speed. The take-off speed is the speed at which the foil begins to fully support the rider's weight and the rider's board is no longer touching the surface of the water. As the rider increases speed at a given pitch angle, the lift from the main foil increases faster than the lift from the stabilizing foil increases as a result of the larger relative area of the main foil. This means that the angle of attack of the main foil must decrease in order to provide the rider with a consistent amount of lift as he or she picks up speed. As the pitch of the kitefoil decreases to reduce the angle of attack of the main foil, the angle of attack of the stabilizing foil increases and there is a relative increase in the stabilizing moment of the foil for a given amount of lift. This phenomenon manifests itself in the stance of the kitefoiler as shown in Figure 3, where the rider centers his or her back foot above the strut of the kitefoil and uses his front foot to apply a force to counteract the stabilizing moment of the kitefoil.

Figure 3. Kitefoil rider stance
http://www.surfertoday.com/images/stories/kitefoilspeed.jpg

The stabilizing moment essentially requires a lift that actively counteracts the lift of the main foil, which decreases the lift-to-drag ratio of the foil and ultimately decreases the foil's efficiency. This is a necessary dynamic effect that ensures the rider's ability to balance on the foil. In simplified terms, the main foil acts as a fulcrum to which the rider applies his or her weight, while balancing the downforce from the stabilizer with his or her front foot. As opposed to balancing on the main foil alone, which could be considered a form of unsteady equilibrium, the stabilizing moment and the rider's need to counteract it contribute to a more stable equilibrium that contributes to the rider's ability to build and maintain speed in variable winds as well as through maneuvers such as tacks and gybes. It is the task of the kitefoil designer to create a design that is able to have a consistent lift force across a wide range of operating speeds while still producing a stabilizing moment that allows for a rider to find stable equilibrium across the same operating range.

COMPUTATIONAL FLUID DYNAMICS

Reynolds-averaged Navier-Stokes (RANS) equations based computational fluids dynamics (CFD) analysis was employed in this analysis as it is able to accurately resolve boundary layer and vorticity effects in the flow. Though a panel code program, such as XFLR5, may be used to analyze inviscid flows with empirical corrections made to account for viscous effects, the complex 3D geometry involved with a kiteboarding hydrofoil, as well as a desire for the most accurate results possible led to the choice of RANS-based CFD for the flow simulation in this thesis.

Turbulence

Turbulence, which can play a significant role in the structure of flow around a body, must be accounted for in CFD analysis. Turbulence is unsteady aperiodic fluid motion in which the velocity components vary and matter, momentum, and energy are mixed (Fast, 2014). Turbulence must be modeled in simulations in order to predict the flow characteristics without being required to compute the unsteady turbulent flow pattern. While the exact solution of the turbulent flow structure may be obtainable, a turbulence model is necessary in CFD to reduce the mesh and run time by eliminating the need to resolve turbulence in the flow. A turbulence model for CFD incorporates a time-averaged

estimate of the fluid conditions in a turbulent region (Fast, 2014).

PARAMETER-BASED OPTIMIZATION

Parameter-based optimization is an efficient means of optimization in which the important features of an object are described within its parametric geometry. As opposed to parameter-free optimization, parameter-based optimization allows for direct variation of the geometric features that are deemed significant to a design's performance.

The type of parametric model used for optimization, known as an "engineering model", attempts to define the significant aspects of the geometry in as few parameters as possible. An engineering model may even purposely omit geometric features that are not considered significant, or disable variation of parameters of little interest, known as nuisance parameters. Parameter-based optimization is split into fully and partially-parametric categories. In the first, the entire geometry to be optimized is defined parametrically, and, in the second, only the features to be optimized are defined parametrically. The optimization procedure outlined in this thesis is of the partially-parametric type, parametrizing only select features of a base-model in an effort to refine its performance.

Partially parametric optimization allows for relative ease in the modelling process but inherently contains less capability than a fully-parametric model. Stefan Harries (2014) notes that "it is more difficult to excite large (game changing) modifications" when optimizing a partially parametric model. For this reason, the partially-parametric optimization procedure is more useful for later-stage design refinement than for early concept design.

MATHEMATICAL OPTIMIZATION

Mathematical Optimization is defined as a technique for finding a maximum or minimum value of a function of several variables subject to a set of constraints. In simple terms, this means finding the best option among available alternative solutions. Any problem that has a solution that can be assessed through an objective function and has one or more variables that contribute to the value of the solution may be subject to mathematical optimization. Figure 4 is an example of a basic optimization, where variables X and Y are changed and the objective function is evaluated on the Z axis. In this example the maximum of the objective function is found at (0,0,4) and represents an optimal solution.

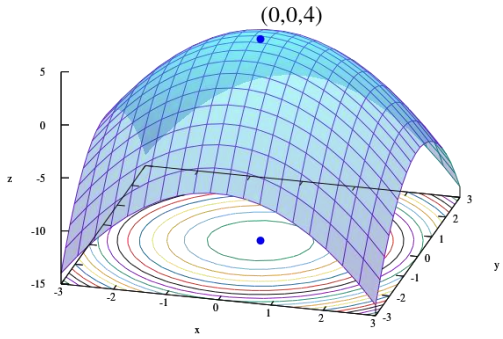

Figure 4. Global maximum on the surface of a paraboloid https://commons.wikimedia.org/wiki/File:Max_paraboloid.svg

The position of a local or global maximum of the objective function may be found where the first derivative of the objective function is zero. This may be solved for mathematically if the objective function is algebraically defined; however, for the optimization of complex problems that require numerical analysis, the objective function may not necessarily have a continuous functional description. In this case, the objective function must be analyzed point by point in an attempt to define its shape in hyperspace with respect to the optimization variables. Through this type of analysis the gradient of the objective function may be investigated in an effort to locate local and global maxima. A global exploration of the design space on a point-wise basis is often used as the first step of an optimization process involving numerical analysis. When a sufficient number of design points have been investigated

to sufficiently map the global trends in the objective function, the investigation of potential local optima can begin through an exploitative, as opposed to explorative, means. An exploitative optimization technique uses the known trends in the objective function to drive variation in the design variables in the direction of the favorable objective function gradient.

Optimizations in which there are multiple objective functions may develop what is known as a Pareto efficiency, which is a design point in which it is impossible to favorably improve the value of one objective function without unfavorably affecting the value of any other considered objective functions. A Pareto efficiency can manifest itself in any multi-objective optimization as a Pareto frontier, in which the points on the frontier are dominant in their optimization of the objective functions as compared to other feasible but ultimately unfavorable designs. Figure 5 shows an example of this, with points A and B lying on the frontier as they optimize, in this case by minimizing, both objective functions f2 and f1 concurrently. The points to the upper right of the red frontier show valid design points that are less favorable in one or both objective functions than those points lying on the Pareto frontier. Any point on a Pareto frontier may be considered an optimum, and ultimately the selection of a preferred design must be considered with respect to factors such as the desired relative degree of optimization of each objective function.

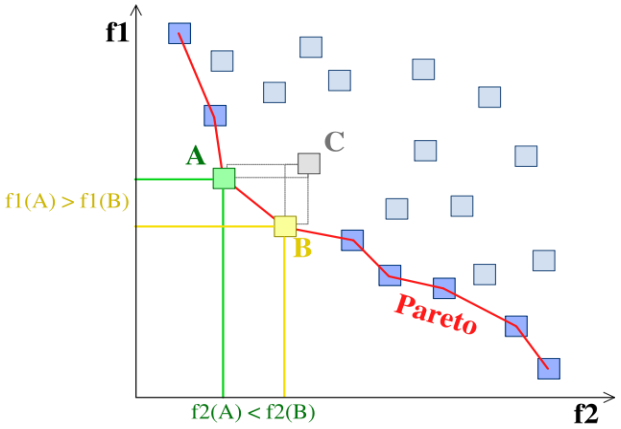

Figure 5. Example Pareto frontier https://commons.wikimedia.org/wiki/File%3AFront_pareto.svg

Figure 6 outlines a generic optimization. First, the design problem is outlined, from which the design variables and objective function(s) are identified. A model to be optimized is then created with the ability to vary the selected design variables, and is then used to explore the problem's design space with respect to the objective function(s). Following a design space exploration, local optima exploitation serves to locate global optima. From this point, a multi-objective optimization problem may investigate the existence of a Pareto frontier with respect to its various objective functions. Ultimately a design optimum is selected from the Pareto investigation in a multi-objective problem or directly from the local optima exploitation in a single objective problem

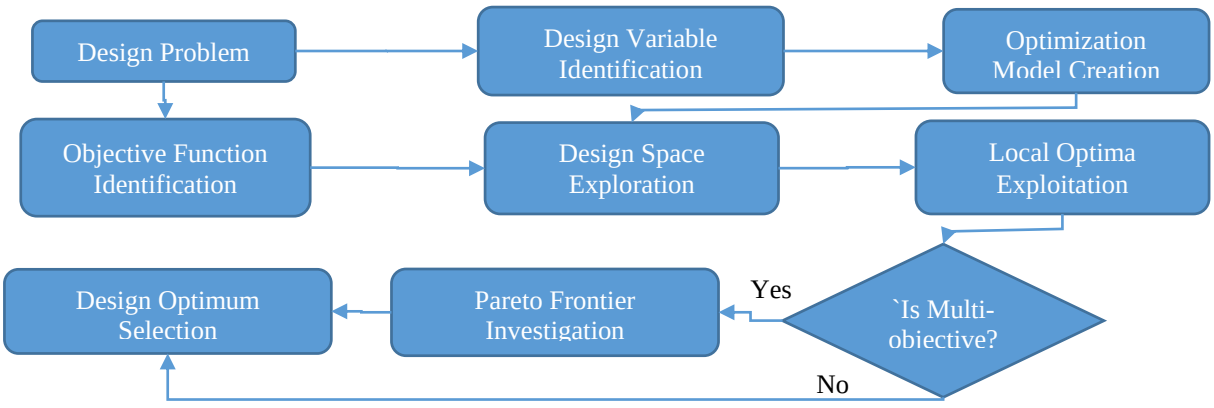

Figure 6. Mathematical optimization sequence

PROCEDURE

DESIGN RESEARCH

In order to determine how best to optimize a foil, it is important to understand the conditions it will operate in and the constraints to which the design is subject. The purpose of the initial research is to determine which foil characteristics are to be varied during the optimization process, determine any geometric constraints that the geometry must follow, define performance constraints that serve to validate design practicality, and to identify an objective function that the optimization will ultimately attempt to minimize.

Design Considerations

The ultimate goal of kitefoil design is to minimize drag and maximize lift while maintaining a foil geometry that is stable enough for a human to ride effectively. While the specific geometry of a hydrofoil relating to an optimum lift-to-drag ratio, or L/D, is dependent upon the operating conditions and constraints of the foil, it is important to note that regardless of the foil operating conditions the hydrofoil must be able to be operated near to the conditions resulting in the optimum L/D in order for the design to be effective. Though a foil design may be sufficient at a given operating condition, the design is not successful if the given operating condition is not attainable by its rider.

Significant Foil Characteristics

The specific geometry of a kitefoil's individual hydrofoils is directly related to the kitefoil's performance. The defining characteristics of the lifting surfaces on a kitefoil are equivalent to those of any three-dimensional airfoil geometry, though in particular the foil section, span wise chord variation, span wise twist, anhedral or dihedral curvature, sweep, and planform are widely varied in currently available designs. The geometry of the strut, fuselage, and the fillet between the two are significant features as well. In order to simplify the optimization process, the unmodified base-model's geometry was used for these components. The significant characteristics of the main foil may be varied along with each other in an attempt to find an optimally efficient main wing. Additionally, in order to produce a constant stabilizing moment, the relative distance between the stabilizer and the main wing along with the size of the stabilizer form a series of acceptable solutions; however, though the varying stabilizer force will require a modification of the lift from the main wing, making this optimization process a three-variable problem.

DESIGN CONSTRAINTS

The design of a kitefoil is subject to a number of constraints that must be considered during the optimization process in order to obtain a practical result. The constraints on a kitefoil generally fall into one of three categories: operational, class rule, and structural/material.

Operational Constraints

Operational limits include constraints imposed by the physical conditions surrounding kitefoil operating conditions. One such constraint resulting from operational considerations is the maximum heel angle. On the water, a larger heel angle will increase the lift produced to windward thus increasing upwind performance. The maximum heel angle is dependent upon the board width, strut length, and foil immersion. The maximum heel angle value is the maximum angle that the kitefoil can be heeled to without the rider's board hitting the free surface. In this analysis the free surface is not considered, and as such the test condition did not consider a heeled foil.

Another constraint imposed under this category is a limitation on the foil span, which if too great may place the wingtip too close to the free surface, resulting in ventilation and ultimately a loss of lift force. This constraint is unnecessary in an "unlimited" optimization process, where the simulations would indicate a loss of performance resulting from such ventilation and thus shy from such geometries in the search of an

optimum. For this analysis, the free surface is not considered; consequently, the geometry is subject to a width constraint in order to effectively supplant the need for ventilation analysis. The width constraint takes the value of the base-case span which serves as a nuisance parameter and is ultimately not optimized for as a free-surface analysis is essential to its effective investigation.

Class Rule Constraints

The International Kitefoil Association, or IKA, specifies in its class rules that the minimum length of the kitefoil is 500 mm measured from its extreme points. There is no maximum length specified. The minimum length is measured perpendicularly from the bottom of the board to the extreme point of the kitefoil. Additionally, the class rule specifies that only a single appendage is allowed and that the appendage shall have the primary purpose of creating vertical lift. The class rule does not impose any limitations on the material of the kitefoil.

Structural/Material Constraints

Structural and material considerations need to be taken into account in order to produce an optimized geometry that may both be constructed accurately from the desired material and be structurally capable of enduring the loads that it is subject to. These considerations impose practical constraints on kitefoil geometry in terms of minimum thickness of the foil sections and strut, as well as maximum allowable stress in the kitefoil. Because the effective implementation of these constraints would require a further analysis of the construction process and potentially a finite-element model coupled with the flow simulation, structural and material constraints are not considered in this optimization.

MODELLING

All parametric modeling was performed in the CAESES-FFW, which is a software that is designed to seamlessly combine CAD modelling, computer simulation, and optimization techniques in a framework, enabling a streamlined design process for refined flow-exposed components. Rhinoceros 3D was used to investigate the base-model and to troubleshoot geometries exported from CAESES.

Base-Model

Rather than creating a kitefoil from scratch, it was determined that a base-model would be best to serve as the starting point for optimization. Using a base-model that has already been built and ridden successfully allows for a realistic base case to be established founded on the estimated operating condition. George Hradil of Delta Hydrofoil provided a CAD model of his company's production hydrofoil (Figure 7) to be used as a base for optimization. As the actual lift force produced by the hydrofoil is unknown at this point, the designer recommended 180 lb of lift as a design point. The CAD model provided was ultimately dissected in Rhinoceros 3D and recreated from its base curves parametrically in CAESES. It is worth noting that the location of the main foil relative to the strut in the vertical (Y) direction was raised slightly to accommodate a wide range of change in main foil angle of attack.

Figure 7. Delta hydrofoil
George Hradil

Foil Design

After determining the exact parameters to be optimized and included in the parametric definition, the design of the model can begin. The

development of the foil starts with the foil section to be used, from which a feature definition is developed. Feature definitions are command sequences written in the native CAESES "Feature Programming Language," which essentially create a functional parametrized definition of the desired foil geometry. The feature definition is then used in a curve engine which takes the feature definition and base section curve as its arguments. The curve engine also requires input for the parameters previously defined in the feature definition. By assigning parameter values to a range specified by an arbitrary function, one is able to create a section with geometry varying according to arbitrary functions over its swept path. The curve engine is used to connect the parameters created in the feature definition to a continuous description of the base curve that varies according to the various function curves. In the later optimization process, it is the "Design Parameters" defining the functions that will be varied rather than the foil itself. In this way, the number of variable parameters is kept to a minimum in order to reduce the computational demand. The curve engine is then incorporated into what, in CAESES, is known as a "metasurface," which is used for creating the complex shapes with deformations and translations that vary along a path defined by a curve (see Figure 8). The variations among the foil shown include anhedral, span wise chord distribution, and span wise twist, applied in varying degrees on each of the three variants.

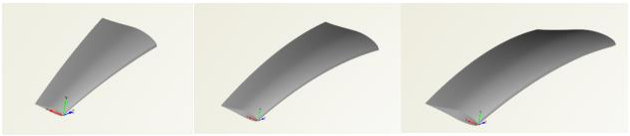

Figure 8. Parametric foil variations in CAESES

Once the foil surface is created, it must be adapted so that it matches the base-model to be optimized. In this way, the base-model geometry is kept constant in the initial case but is programmed with the geometric variability that is deemed significant to the foil's performance. The plan and profile view curves of the leading edges of both the front and rear foils was converted into 2D in Rhino and then imported as an IGES file into CAESES where the curves serve as the function curves defining the foil geometry. As in the base-model, the trailing edge is straight in the plan view. Ultimately, the size of the stabilizer, the main foil angle of attack, and the fuselage length serve as design variables in the optimization process.

Strut/Fuselage

For the sake of simplicity, the base-model strut, fuselage, and transition between the two is kept as a constant geometry in the optimization process. The geometry from the base-model was sectioned in Rhino and lofted in CAESES to form a Brep surface. In order to investigate the tradeoff between fuselage length and stabilizer area, the length of the fuselage must be parametrically variable.

To adjust the length of the fuselage, a transformation function within CAESES known as a "Delta Shift" transformation was employed, which enabled the fuselage to be effectively expanded or contracted in the *X* direction, lengthening or shortening the portion of fuselage aft of the strut. The fuselage was lengthened only aft of the strut, so that the only effective change was the distance between the main foil and the stabilizer, with no relative change between the main foil and the strut taking place. The Delta Shift transformation applies a scale factor according to a curve known as a "Delta Curve," which is defined in two dimensions and specifies the magnitude of the scale factor relative to the position along the length of the fuselage. The Delta Shift was used because (1) it allowed for a continuous description of the scale factor to be applied to the fuselage which enabled the fuselage surface to remain fair, and (2) it allowed for translation to be applied solely in the fuselage region aft of the strut. Had a simple translation been used, the fuselage would not have maintained its validity as a continuous surface, and a discontinuity would have occurred from the

point where the translation was applied. The Delta Curve ranged from 0 to 1 on the abscissa. The ordinate of the Delta Curve was zero until an abscissa value of 0.31, which is the percentage of the length along the fuselage that the strut intersected the fuselage, after which the ordinate continuously varied from 0 to the desired fuselage extension factor. In this way the fuselage may be lengthened or shortened while maintaining a fair transition between the original and scaled regions of the fuselage geometry. The fuselage extension factor, which defines the extension or shortening of the fuselage, serves as a design variable. A Boolean operation was used on the fuselage and strut to join them into a single surface.

DESIGN PARAMETERS

The final design parameters chosen are detailed in Figures 11-12, with the base-case geometry, an exact recreation of the delta hydrofoil base model, shown in Figure 9 for comparison. It is worth noting that the magnitude of the parameter variations in the examples is relatively large to exaggerate the change in geometry and to make clear what component of the design is being altered. The design parameters were selected specifically to investigate the performance tradeoffs between different geometric variations while also being able to produce a foil that created the required lift force and moment.

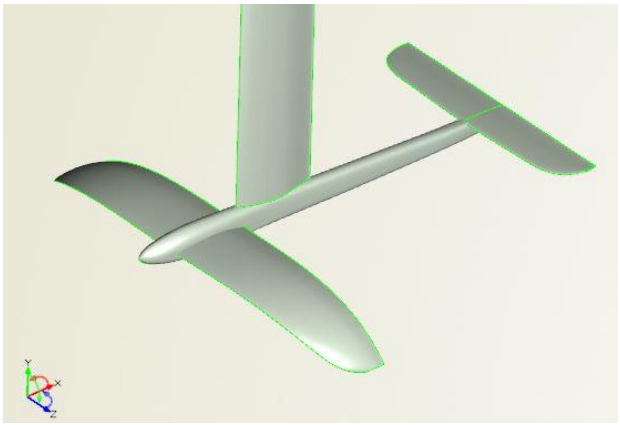

Figure 9. Base-case parametric geometry

Change in Main Foil Angle of Attack

The change in the angle of attack of the main foil relative to the base case is known as the change in main foil angle of attack, $\Delta\theta$. The change in main foil angle of attack is effectively a rotation of the main foil about the Z axis with the origin located at its leading edge. An increase in this parameter, as shown in Figure 10, generally increases the main foil's lift. In the optimization of the stabilizer, the change in main foil angle of attack serves to offset the change in lift due to the variation in stabilizer down force. The sign of this parameter is negative when angle of attack is increased.

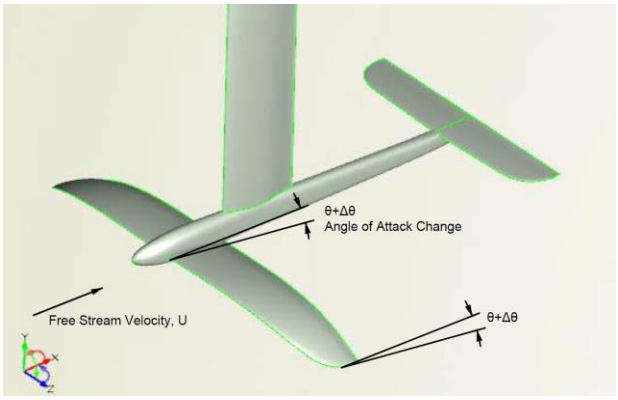

Figure 10. Change in main foil angle of attack

Fuselage length scale

The scaling of the fuselage in the *X* direction is known as the fuselage length scale. The measurement of the extension is designated as ΔL in Figure 11, while L designates the base length of the fuselage. The parameter serves to alter the distance from the stabilizer foil to the main foil. The scaling is applied to the region of the fuselage aft of the trailing edge of the strut. The mechanism used to apply the scale factor was a Delta Shift transformation in CAESES, as discussed previously.

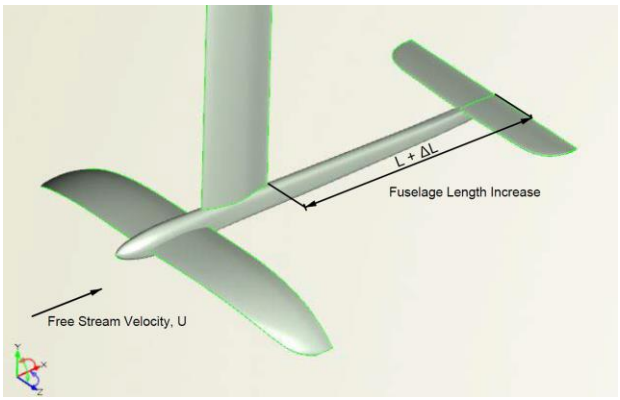

Figure 11. Fuselage length scale

Stabilizer Foil Scale

The 3D scaling of the area of the rear foil relative to its trailing edge is known as the stabilizer foil scale. An increase in this parameter serves to increase the downforce created by the stabilizer. The effect of an increase in the main foil scale is shown in Figure 12.

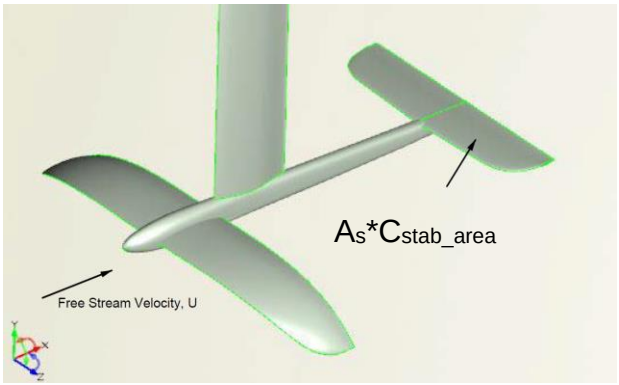

Figure 12. Stabilizer foil scale

Parameter Interdependence

The optimization of the stabilizer involved variation of the fuselage length scale, the stabilizer foil scale, and the main foil angle of attack. While each design variable may be altered independently to investigate its impact on the kitefoil performance, correlations between the design variables that inherently account for the impact of the change of one on the others may help to further refine the optimization, while considering the optimization of the fuselage length, stabilizer area, and main foil angle of attack. As the main foil angle of attack decreases, the lift produced by it decreases, which in turn decreases the stabilizing moment. Because the total lift produced decreases, the stabilizer area must decrease. The stabilizer area must decrease because it must produce less downward lift to balance the main foil lift decrease so that the design lift may be obtained. As the stabilizer lift decreases, so does the stabilizing moment produced by it, and, as such, the fuselage length which creates the moment arm between the two lifting surfaces must increase in order to return the stabilizing moment to its design value. Ultimately, as the main wing angle of attack decreases, the stabilizer area must decrease and the fuselage length must increase.

It is clear that all valid design variants must follow these correlations, and, in order to improve the efficiency of the optimization process, it is worth considering them in the definitions of the design variables themselves. The necessary positive correlations between these design variables means that, instead of independently defining design variables in the optimization process, variable gain factors can be used to identify the correlated design variables relative to a chosen base design variable. To represent these correlations in the variation in geometry, the fuselage extension factor was the result of a positive gain factor multiplied by the change in main foil angle of attack, and the change in stabilizer area was the result of a negative gain factor multiplied by the change in main foil angle of attack. These stabilizer area and fuselage extension gain factors were then selected as the design variables for the actual optimization process in addition to the reference design variable of change in main foil angle of attack.

COMPUTATIONAL FLUID DYNAMICS

Geometry

The geometry used for the hydrofoil surfaces was imported from CAESES in .STL format as shown in Figure 13. The mesh resolution of the CAESES stereolithography export was sufficient to accurately resolve the forces and moments in the flow; however, in the surface pressure

representation, there were unphysical gradients that resulted from the geometry approximation that the .STL mesh distribution produced. The unphysical gradients can be identified as the span-wise bands of nearly constant pressure, which correspond to seams in the .STL mesh.

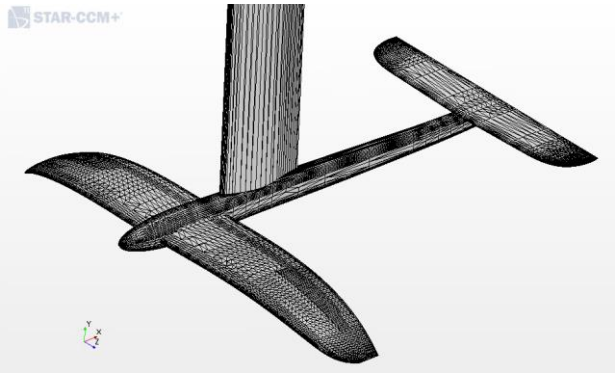

Figure 13. Imported .STL geometry

Fluid Domain

The computational domain that defined the boundaries of the flow simulation is shown in Figure 14, with the direction of travel aligned with the *–X* axis and the port and top domain boundaries hidden. The red surfaces, along with the port and top boundaries, are velocity inlets, with 20 knots of saltwater flowing in the positive *X* direction. This condition represents the kitefoil operating in undisturbed open water, with the effects of free surface being neglected. The orange boundary is the pressure outlet. The blue boundary is a symmetry plane which bisects the flow domain through the centerline of the kitefoil. Because the foil test operating condition does not have an angle of heel or pitch, a symmetry plane about the *XY* plane was used to simplify the model and halve the number of cells in the computation.

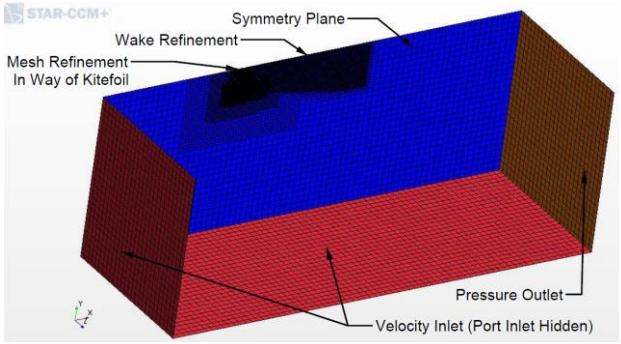

Figure 14: CFD flow domain

The box that encompasses the domain was created with a "block" shape part in Star-CCM+. The spacing of the domain boundaries from the hydrofoil were sized iteratively, increasing the size of the domain until the pressure distribution on each face was nearly constant.

Fluid Volume Meshing

The final flow domain surfaces were created through the use of a Boolean subtract operation between the hydrofoil surface and domain boundaries. The Boolean subtract operation trims surfaces where they intersect other surfaces, such as at the intersection of the main wing and the fuselage at the root. The surfaces were not joined so that the individual parts could later be identified for mesh refinements in CFD. An automated mesh operation was defined in Star-CCM+ to create the fluid volume mesh for analysis.

Prism layers are used to transition the mesh from the surface geometry to the free-stream geometry and allow for accurate simulation of boundary layer flow. A detail view of the prism layer surrounding the main wing is shown in Figure 15. It is worth noting that 14 prism layers were used, with a prism layer stretching factor of 1.706 and a prism layer total thickness of .002833 m, to result in a Y+ value below 1 at the surface. The Y+ value is a Reynolds number based on the cell closest to the foil surface, and is an indication of the degree of boundary layer refinement, with lower values corresponding to more refinement.

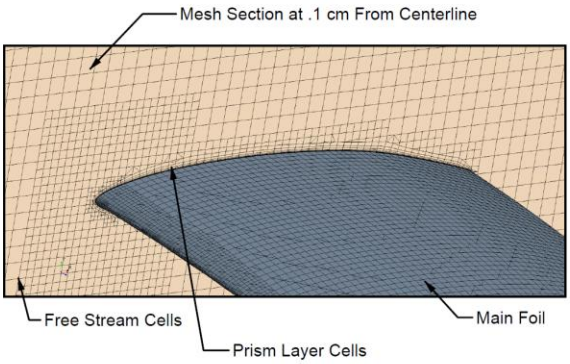

Figure 15. Prism layer refinement on main foil

What are known within the Star-CCM+ automated mesh operation as "Custom controls" were used to refine the volume mesh in an attempt to efficiently and accurately simulate the flow around the foil. A "surface control" is a type of custom control that applies a user-defined mesh refinement to the cells surrounding a selected surface. One such surface control was applied at the domain boundaries to the disable the prism layer and specify a target surface size that is relative to the base cell size. This control effectively simplified the mesh near the domain boundaries where the flow is nearly uniform. A surface control was also used on the main foil and stabilizing foil surfaces to specify the target surface size and minimum surface size, both of which were relative to the base. This surface control was used to refine the mesh of the lifting foil surfaces in order to more accurately capture their shape in the finite element mesh. Finally, a wake refinement was applied to the entire hydrofoil geometry which gave the volume mesh additional resolution in the region of disturbed flow behind the foil. The wake region created had a distance of 2.5 m and a spread angle of 10 degrees.

Physics Conditions

A turbulence model is necessary to accurately simulate steady turbulent flow because of the highly chaotic and unsteady nature of actual turbulence. The eddy-viscosity-based K-Epsilon Two-Layer turbulence model was selected for its accurate turbulence simulation of well-streamlined bodies. Additionally, a segregated flow model was employed because it is a suitable approximation to make in the design operating condition, where the temperature, pressure, and density of the fluid vary only slightly. The reference pressure is atmospheric, and the flow velocity is 20 knots.

Simulation Validation

The base-case flow simulation was run for 250 iterations, after which the resultant forces were measured to differ by roughly .009% between iterations, and the residuals dropped to values below $5*10^{-3}$. The final surface pressure distribution is shown in Figure 16.

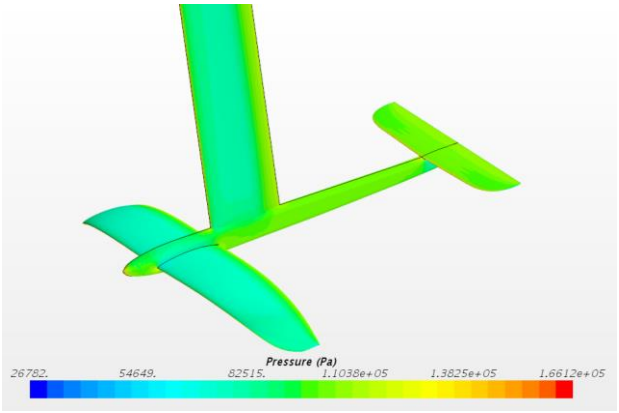

Figure 16. CFD base-case surface pressure

The lift of the main wing and stabilizer, the total lift, and the total moment were validated using a first-principles analysis. Using XFLR5, the lift coefficients of the main and stabilizer foil sections was determined at the design condition.

Corrections for the 3D effects on lift coefficients were determined through Prandtl's and Helmbold's formulae. A summary of the first-principles kitefoil analysis compared with the CFD results is shown in Table 1.

Table 1. CFD validation summary

	Units	1st Principles	CFD	% Error
Main Foil Lift	lbf	296.91	260.1	12.41%
Stabilizer Lift	Lbf	88.617	88.7	-0.12%
Total Lift	Lbf	208.29	178.8	13.58%
Total Drag	Lbf	28.56	27.5	-3.71%
Total Moment	Lbf*ft	80.06	75.6	5.52%

The CFD results differ from the first-principles estimate largely because of the effects of viscous three-dimensional flow. In particular, the discrepancies are caused by the presence of turbulence and finite-span effects. Turbulence alters the flow field around the foils and also causes a loss of energy in the flow. Finite span effects include the loss of energy in the flow to wing tip vortex generation. The estimate of the stabilizer lift is very accurate because the wing has a very rectangular planform with no dihedral or anhedral, which lends itself well to 3D lift coefficient estimations. The lift produced by the main foil is comparatively lower than the first principles analysis because of the variation in planform of the main foil as well as its anhderal curvature.

OPTIMIZATION

CAESES provides a framework for the optimization of parametric geometries modelled within the program. Design variables from the modelling process are varied in design engines, which modify the variables in accordance with either an exploratory or exploitative method.

Automated Geometry Variation

The five design variables serve adequately to vary the performance-critical geometry. With these variables, the engineering model is equipped with the necessary geometric functionality to accommodate significant design variation that may increase performance. CAESES design engines allow for the automated variation of design variables subject to user-defined ranges.

Geometric limits are placed on the parametric model in the form of maximum and minimum values for each of the design variables. The ranges of the design variables were explored through a manual analysis of the parametric model. The design variable limits were applied to keep the final export geometry valid and acceptable for CFD simulation, which included the absence of any modelling errors such as open edges or singularities. The limits on the design variables also took into consideration that practical limits of kitefoil geometry in terms of its physical construction and integrity, though ultimately no external geometric constraints related to construction or strength were imposed.

Result Constraints

The design engines implemented in the optimization procedure are subject to performance-based constraints that pertain to the vertical lift and stabilizing moment created by the hydrofoil. The base-case under the design operating condition provided the design lift and stabilizing moment, which are considered necessary for the steady-state operation of the foil under kite power. Ideally, each variation in geometry would be tuned well enough to provide the exact same lift and moment as the base case, though in reality this would require an excessive number of design iterations. This would limit the effectiveness of the optimization as a whole, given that the overall computational effort is limited by time. Instead, the variations in geometry that resulted in lift within 5% of the base case and moments within 10% of the base case were considered valid designs. This criteria is implemented to reduce the computational

effort of the optimization, and, while it is true that a smaller percentage would create a more accurate optimization, for the purposes of this study a reasonably large range of acceptable lift and moment was used to maximize the potential of finding an optimum with a reasonable amount of computational effort. The range on lift is relatively small because no matter the operating condition a kitefoil must produce a nearly constant lift to offset the weight of the rider. The range on moment is comparatively large because the required stabilizing moment is not necessarily a single value, but rather is dependent upon the preference of the rider. As such, it is assumed that while the base case stabilizing moment is an appropriate value for normal kitefoil operation, if the foil were to produce a moment within 10% of the base case, the foil would still be able to be ridden efficiently with an accompanying adjustment of the rider's technique.

The lift validity criterion is based on the assumption that a foil that produces a magnitude of lift close to the design lift with an improved L/D would still operate with an improved L/D when it is operating with a pitch angle that allows it to produce exactly the design lift. This assumption can be validated by running the optimized foil through an additional iterative process in which foil pitch alone is adjusted until the foil produces the design lift; however, it is clear that as a result of changing the pitch angle of the foil would likely modify the resultant moment. Ultimately, a foil optimized with a range of acceptable performance may never actually be able to operate at the original design condition, even with a change in foil pitch. The desired lift of the foil is a constant that depends upon the rider weight and kite force, while the desired stabilizing moment is subject to the rider's preference. As such, minor changes in the stabilizing moment of the optimized foil that produces the design lift would be considered acceptable, though in practice the design may be less preferable to a rider. These validity criteria are established to decrease the computational demand required to find an optimum.

CAESES/STAR-CCM+ coupling

In order to set up CAESES for optimization, it is necessary to couple the design engine, which drives the selection of geometry variants and performs the results analysis, to the flow simulation, which provides the performance results. When properly coupled, CAESES will run a CFD simulation for a geometry variation and import the result files automatically. In an optimization process, CAESES automatically uses the coupled simulation software to evaluate design iterations as soon as they are created. This is made possible through the use of the CAESES software connector interface. The CAESES software connector has five main parts: input geometry, input files, result values, result files, and the Runner. To define the geometry used for each iteration, the input geometry is selected from the CAD generated in CAESES, the file name and type are defined, and the final export location is defined with a path that is relative to the local optimization directory.

JavaScript macros are used to communicate commands from CAESES to Star-CCM+ and serve as the input files. The JavaScript macros are recorded in Star-CCM+ and modified in CAESES by incorporating variable parameters, such as the desired number of simulation iterations, and relative file locations, which allows Star-CCM+ to locate the geometry to import. Five JavaScript macros were created for this coupling process: "part replacement", "meshing", "start simulation", "outputs", and "master macro." The part replacement macro replaces the base case geometry with the geometry variant to be analyzed. The meshing macro executes the Boolean operation between the new geometry and the flow domain, and then executes the automated mesh operation. The start simulation macro defines the number of iterations for which the simulation will run and executes the simulation. The outputs macro saves the .csv files of the resultant forces and moments,

as well as the simulation .case file to a relative location that is specific to each design iteration. The master macro runs all of the previous macros in the necessary order.

The Runner executes the Star-CCM+ program according to a provided configuration. The configuration is the database that provides arguments to the input macros that have previously been chosen in CAESES. For instance, when recording the macro to save a result file, an absolute file reference must be chosen in Star-CCM+. The absolute file reference will still be in the macro once it is imported into CAESES, though in order for the result data of individual iterations to be saved independently, a relative reference to the design directory must be created. This is done by selecting the file reference of the macro while in the CAESES software connector macro editor and redefining it with a string parameter. The string parameter then shows up as a value in the configuration, where it may be defined as a reference to the relative location of the design directory. The Runner also requires arguments which it provides Star-CCM+ with upon start up. The required arguments include a reference to the license server, the number of processes to run, the command to run in batch mode, the relative location of the master macro JavaScript file, and the absolute location of the base simulation file to open. The specification of batch mode is required because it runs the software without the graphical user interface, which enables CAESES to upload the result files automatically. Ultimately, the batch mode operation puts all CFD analysis involved with the optimization behind the scenes, leaving only the resultant simulation data behind in the design engine results pool. The simulation application itself must be specified within the Runner in the form of a "Resource Manager", which requires the absolute path of the Star-CCM+ executable file as well as the maximum number of running instances. Many optimization engines are able to run more than a single simulation at a time, and, as such, four running instances with ten processes each were specified to efficiently compute the simulations.

After CAESES has run a simulation and the result files have been saved in their correct locations, the result values are imported based upon a sample result file. CAESES allows for the definition of parameters based on the relative selection of data within the sample .csv file. To import the results for lift, drag, and moment into CAESES the column which corresponds to the desired measurement of the .csv file is chosen and the last row of the file is selected in order to import the results of the last completed iteration. The imported result values are used as evaluation parameters to drive the optimization process. Similarly, result files can be imported and analyzed within CAESES for visualization of the simulation results for each iteration.

Design Engines

Design engines are optimization tools that drive the variation of the design parameters in ways that effectively investigate the design space. Design engines require the specification of the design variables to be modified as well as a range of values for the design variables to be modified over. Design engines also require the specification of evaluations, which are parameters that may be related to the geometry of the design iteration or related to the simulation results. An evaluation may be marked as an objective in an exploitative design engine, in which case the objective is sought to be minimized; for this case the L/D objective function was multiplied by negative one. Additionally, design engines allow for the designation of constraints to follow which determine the validity of each design iteration. Design engines serve to either explore the design space or to exploit local minima or trends of performance. The design engines used in this optimization were the Sobol and T-Search.

Design space exploration is accomplished through the use of Design-of-Experiments (DoE). Design-of-Experiments are mathematical algorithms that efficiently explore multi-variable

problems with minimal computational effort. DoE enable the user to get a better understanding of what parameters in their optimization are significant and what limits should be placed on their variation. Design space explorations can also be used as a starting database for exploitative design engines to analyze and assess trends between the inputs and outputs. The Sobol design engine provided within CAESES was used for exploration in this study. The Sobol is a deterministic algorithm that produces a seemingly random distribution of the design variables based on a given number of variants. The Sobol engine selects where to place its next design variant in a way that systematically interrupts the largest regions of unexplored design space.

With the design space sufficiently explored, the investigation of local objective function minima can begin with exploitative design engines. An exploitative design engine uses a mathematical algorithm to drive design variation in the direction of objective function minima based upon trends in the design space. A Tangential Search, or T-Search, is a design engine that investigates the design space surrounding a local minimum and produces new variants based on the relationship between the design variables and the descent of the objective function. A T-Search essential analyzes the gradient of the design variables in hyperspace with respect to the objective function. A T-Search is unable to identify any optima as local or global, and it is important that the design space is thoroughly explored prior to a T-Search in order to identify potential global trends.

OPTIMIZATION WORKFLOW

The procedure followed to obtain an optimized foil geometry involved significant refinement and definition of the design space in order to obtain meaningful results within a reasonable number of computations. With an unlimited amount of time, the entire design space could be exhaustively explored, and an optimum could be found through brute force alone; however, this optimization procedure was subject to time constraints, and, consequently a significant effort was made to improve the efficacy of the design variation. The method of design variation was refined through the identification of correlations between the design variables that produced valid kitefoil geometries so as to reduce the number of invalid geometry.

Initial Exploration

An initial exploration of the design space was performed using a Sobol design engine with 200 iterations, in which the values of the design variables were varied according to the limits in Table 2.

Table 2. Initial Sobol Design Space Exploration Design Variables Ranges

Design Variable	**Lower Limit**	**Upper Limit**
Main Angle of Attack	-.8	.8
Fuselage Extension	-.2	.6
Stabilizer Area Change	.6	1.3

Ultimately, this preliminary exploratory design space survey was completely unfruitful, as the design space was simply too large to effectively explore. The modification of the design variables independent of one another proved to be highly inefficient because, as discussed in the parameter interdependence section, the Sobol design engine allowed for relative variation in the design parameters that would inherently not produce valid geometries. The ineffectiveness of this computationally intensive study lead to the implementation of gain factors to relate the design variables fuselage scale and stabilizer area to the main wing angle of attack change.

Subsequent design space exploration was performed with the parameters of fuselage extension and stabilizer area scale replaced with parameters called fuselage extension gain and stabilizer area gain. Starting with a large range

for both gain factors, the main wing angle of attack was fixed at 0.3 degrees, and a Sobol exploration was performed for 200 iterations. Because the lift coefficient curves for the main wing is nearly linear for small angles of attack, the isolation of the change in main foil angle of attack theoretically still allowed for a stabilizer area gain value to be found that would apply over a range of main wing angles of attack. In theory, the length of the fuselage should have little to no impact on the total lift produced by the foil, so the stabilizer area gain value was found where the lift produced was closest to the design point. The stabilizer area gain value was found to be roughly between -.17 and -.22, shown in Figure 17 as the cluster of blue points that settles around a lift/2 value of 90. For all CAESES generated graphs, blue dots represent valid geometries while blue X's represent geometries that failed the lift or moment constraints.

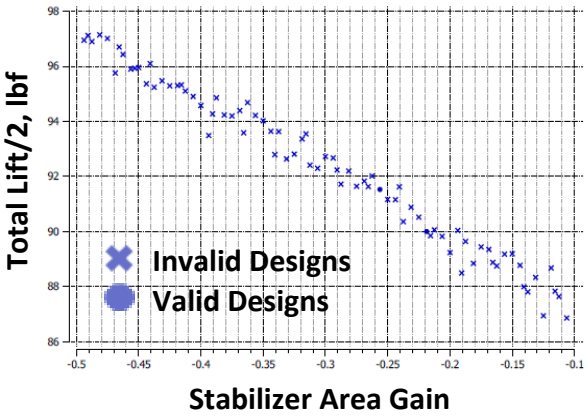

Figure 17. Stabilizer area gain vs. lift/2

A point of stabilizer and fuselage gain values that satisfied the lift and moment constraints was found at the stabilizer area gain of -0.22 and a fuselage extension gain of 0.535.

These gain values were then investigated over a range of angles of attack and expanded to a limited range of -.24 to -.21 for the stabilizer area gain and .52 to .55 for the fuselage extension gain. The main wing angle of attack change was varied from -0.1 to 0.5. Figure 18 shows a positive and linear correlation between decreasing the main wing angle of attack and increase in the difference between the design lift and the result lift (the lift difference is an absolute value, the actual sign of which is negative up to a change in main foil angle of attack of approximately 0.1). This study provided evidence that the stabilizer area is linearly correlated to the angle of attack of the main wing and supports the theory that an accurate gain value would allow for the stabilizer lift to balance the main foil lift, producing a constant total lift force, across a range of changes in main foil angle of attack. While this study validated the theory behind the implementation of gain factors, it was not thorough enough to refine the values of the factors themselves.

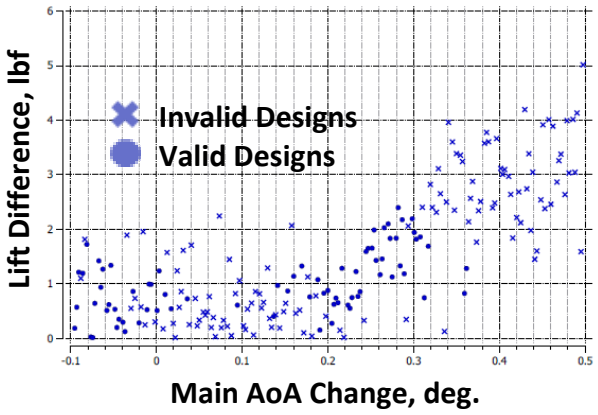

Figure 18. Main wing angle of attack change vs. lift difference

Model Refinement

The stabilizer area gain was first investigated in detail because it has an effect on both the stabilizing moment and the total foil lift force. By solely considering its impact on the total foil lift force and neglecting the stabilizing moment, the stabilizer area gain may be determined, leaving the moment correction to the fuselage extension gain, which has no impact on the foil total lift. To further investigate the stabilizer gain value, a Sobol exploration in which the fuselage extension gain was set at a constant while the stabilizer area gain was varied from -.26 to -.25 was performed. Ultimately a stabilizer gain of -

.255 was identified as an acceptable starting point.

With a value for the stabilizer area gain determined, the fuselage extension gain could be investigated by analyzing the moment across a range of main wing angle of attack changes. Initial attempts at identifying a constant fuselage extension gain value were unfruitful, and, through a process of trial and error, a polynomial relationship between the fuselage extension factor and the change in main wing angle of attack was developed. A linear correlation between the change in main foil angle of attack and the fuselage extension gain was determined, and, as such, it was determined that the fuselage extension factor corresponded to the change in main wing angle of attack through a second-order polynomial relationship as defined in Equation 1.

$$C_{fuse\_extension} = \left(\frac{\Delta_{AoA\_main}}{G_{fuse\_extension}} + 0.85 \right) \cdot \Delta_{AoA\_main} \quad (1)$$

The previously determined stabilizer area gain value of -0.255 and the polynomial relationship defining the fuselage extension factor together form an approximate model of valid kitefoil geometry for a range of change in main foil angle of attacks. A simple test matrix from -.5 to 2 degrees of change in main foil angle of attack was executed. Figures 19 and 20 show that the linear model created to approximate kitefoil geometry that produced lift and moment close to the base case started to fail at higher angles of attack, as shown by the invalid designs created past an AoA of about 0.7. Ultimately, an optima was not identified from this model as the gradient of the L/D could not be sufficiently investigated as a result of the constraints on geometry imposed by the linear relationships established.

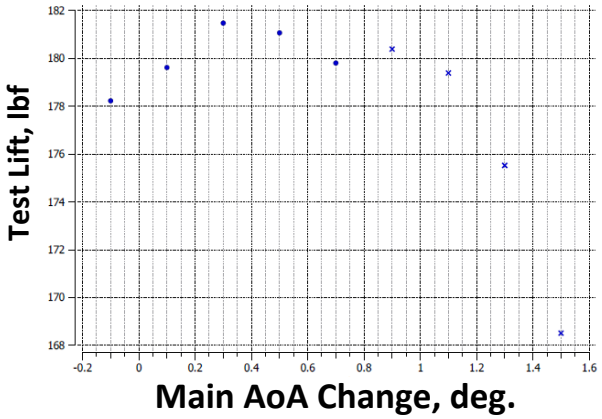

Figure 19. Change in main foil angle of attack vs. lift

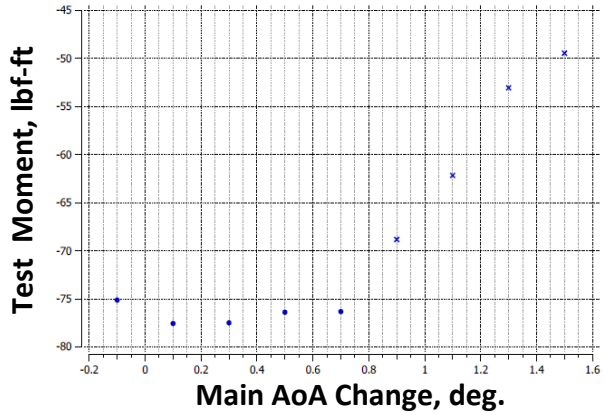

Figure 20. Change in main foil angle of attack vs. moment

A large Sobol design engine was run to explore valid geometries for a larger range of main wing angle of attack changes, specifically from -.2 degrees to 2.2 degrees. The stabilizer area gain value was varied from -.35 to -.245, and the fuselage extension gain, which served as the second-order coefficient in the polynomial expression for the fuselage extension factor, was varied from 3 to 4.2.

RESULTS

Figure 21 displays the results from the large Sobol exploratory optimization of the constrained model and appears to show a minima of the objective function, L/D, around a main angle of attack change of 0.8 degrees. The

minima point is circled in red. Generally speaking, the efficiency of the kitefoil may increase as the main wing angle of attack is relieved, the fuselage extended, and the stabilizer area decreased. It may be possible that at some point in the fuselage length tradeoff the increased frictional drag caused by the extended fuselage would overcome the decreased resistance that came from using a smaller stabilizer foil. The optima found potentially represents a point near to this theoretical L/D maxima.

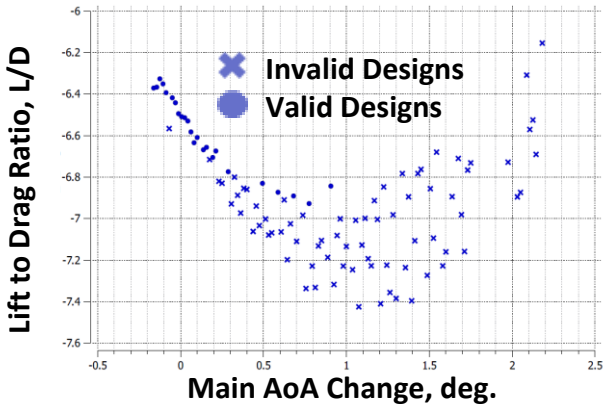

Figure 21. Main wing area change vs. L/D

While an increase in performance was found, this solution is not necessarily a global optima because the gradient in the design space surrounding it has not been exhaustively explored. Figure 22 shows a comparison between the optimized and the base case hydrofoil. The combination of a longer fuselage with a smaller stabilizer wing ultimately produces a more efficient foil for an operating condition equivalent to that of the base case.

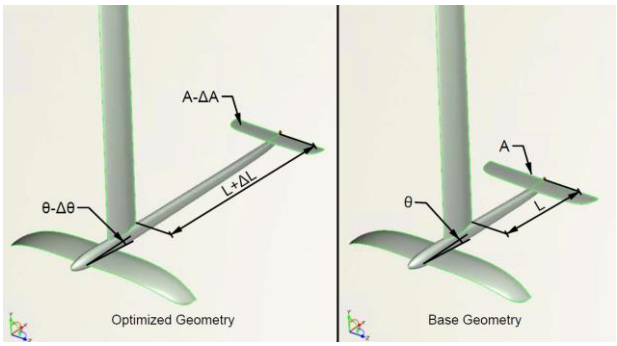

Figure 22. Optimized vs. base case geometry

Table 3 shows a comparison between the most efficient kitefoil found in the optimization process versus the base case.

Table 3. Optimized hydrofoil comparison

	Units	Delta Hydrofoil	Optimized Hydrofoil	% Difference
L/D	-	6.49	6.97	7.29%
Lift	lbf	178.82	181.71	1.61%
Drag	lbf	27.53	26.08	-5.29%
Moment	lbf*ft	75.64	71.97	-4.84%
Fuselage Length	ft	1.82	2.91	59.41%

CONCLUSIONS

Through the creation of a parametric engineering model that encompasses the characteristics of geometry that are significant to performance, it is possible to improve the performance of a given model through an optimization procedure. Depending upon the number and range of variables and constraints to which the model is subject, the scope of the problem in terms of computational demand may be too excessive for practical use. Refinement of the range of geometry variation with respect to the known regions of interest in the design space may significantly improve the effectiveness of the optimization relative to the amount of iterations required; however, it is essential that a model not remain unnecessarily constrained, in which case it may not be able to realize optimal configurations. While a true optimization may result in the discovery of a global maxima of performance, optimization strategies can be used

to improve the performance of a model without ever discovering a true optimum.

The local optimum found in this thesis is in accordance with the theoretical existence of such an optimum, and its identification through the use of a parametric optimization procedure validates the design of the engineering model as well as the feasibility of design optimization in CAESES-FFW coupled to Star-CCM+. Though the operating condition of the optimized foil varies slightly from that of the base foil, the differences may be considered negligible when the large variations of practical riding conditions are considered. The increase in the L/D of the optimized foil, while maintaining a nearly equivalent operating condition as the base case foil, indicates that a hydrofoil geometry subject to a design condition can potentially gain an increase in performance through precise and deliberate change of the design variables that result from a parametric optimization procedure.

REFERENCES

Amelia Nunn (2011). Investigation into the efficiency of a T foil. Australian Maritime College.

Fast, Nathan S. (2014). *A computational fluid dynamics analysis of vertical-lift producing daggerboards for high-performance sailing yachts*. Senior Thesis, Webb Institute.

Harries, Stefan (2014) *Practical Shape Optimization Using CFD*. Friendship Systems. DNV GL.

Hydrofoiling History. (n.d.). Retrieved December 5, 2015, from http://www.hydrofoil.org/history.html

Muller, Christoff. "Design." KiteHydrofoil.com. Cape Town Foilboards, 24 Dec. 2013. Web. 10 Apr. 2015. <http://kitehydrofoil.com/index.php/design/design.html>.

Munson, B. R., *et al*. Fundamentals of Fluid Mechanics (5th ed.), New York: John Wiley & Sons, 2006.

Redondo, G. (2015, October 19). Kitefoil Optimization [Online interview].

Schacht, Roxanne R. and Zangre, Douglas C. (2013). *A computational fluid dynamics analysis of curved daggerboards on high-performance sailing yachts*. Senior Thesis, Webb Institute.

Valenza, D. (2015, October 1). Kitefoil Construction [Online interview].

ACKNOWLEDGMENTS

Numerous individuals have been instrumental to the successful completion of this thesis. I would like to thank my principal advisor, Professor Adrian Onas, for his support and knowledge throughout the project. I would also like to thank George Hradil of Delta Hydrofoil for providing me with the initial CAD model as well as some guidance pertaining to kitefoil geometry. Additionally, I am very grateful to Scott Graham ’83 for his advice and counsel on kitefoil design considerations. I would like to thank Bryan Lake and Marvin Baumeister as well for their input regarding modern racing kitefoil design trends. Lastly, I would like to thank Mr. William H. Webb for his enduring vision of a spectacular educational institution.

Zachary P. Backas completed his bachelors degree at Webb Institute in June of 2016 and is now a naval architect/marine engineer at Shipwright LLC in Ft. Lauderdale, FL.