An Efficient Methodology for Using a Multi-Objective Evolutionary Algorithm for Winglet Design

Frank Kody
Graduate Research Assistant
Department of Aerospace Engineering
The Pennsylvania State University
University Park, Pennsylvania, USA
Fkody1@gmail.com

Goetz Bramesfeld
Assistant Professor
Department of Aerospace Engineering
Ryerson University
Toronto, Ontario, Canada
bramesfeld@ryerson.ca

Sven Schmitz
Assistant Professor
Department of Aerospace Engineering
The Pennsylvania State University
University Park, Pennsylvania
sus52@engr.psu.edu

Abstract

Winglets are designed for the Janus B sailplane through the coupling of an aircraft design code that uses a high-order potential-flow solver with single-objective and multi-objective evolutionary algorithm optimization methods. The use of the single-objective optimizer, Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), serves as a stepping stone to employing a multi-objective optimizer, the epsilon-dominance Multi-Objective Evolutionary Algorithm (ε-MOEA). The multi-objective evolutionary algorithm proves to be successful in designing several winglets with favorable changes in performance. For example, one winglet generated in a two-objective study was able to achieve a 0.1% cruise drag reduction and a 4.5% thermal drag reduction; this results in a peak cross-country speed improvement of 5.5% during weak weather conditions with maximum thermal core strengths of 2 m/s and 0.6% during strong weather conditions with maximum core strengths of 8 m/s. While the design methodology is far from being fully matured, it provides a solid foundation for future research.

Nomenclature

AoA	=	Angle of Attack [deg.]
b	=	span [m]
C_{DP}	=	profile drag coefficient average over span
c_l	=	two-dimensional lift coefficient
C_L	=	three-dimensional lift coefficient
c	=	chord [m]
D	=	drag [N]
$E(t)$	=	archive population
e	=	span efficiency
f	=	objective function
h	=	winglet height [m]
$\Delta K(h)$	=	induced-drag factor dependent on height of winglet
M	=	Mach number
P	=	solution
$P(t)$	=	traditional evolutionary algorithm population
S	=	planform area [m^2]
V_{CR}	=	crossover velocity [m/s]
W	=	weight [N]
ρ	=	density [kg/m^3]
ε	=	epsilon-dominance, hyper-volume dimension

Introduction

The addition of winglets to high-performance sailplanes has become common practice within the soaring community. Examples are the Schempp-Hirth Discus 2 or the Ventus 2. In order to have performance enhancing winglets, these have to be tailored for each particular type of sailplane due to the complex flow conditions that exist near a wingtip. An additional complication is the relatively challenging and diverse mission profile typical for sailplanes that operate over a wide range of lift coefficients at high- and low-speed flight during inter-thermal cruise as well as thermals for altitude gain, respectively. This requires a relatively elaborate design process that generally relies heavily on the personal experience of a designer. If designed incorrectly, a winglet can easily lead to a performance penalty rather than an improvement.

Brief Overview of Winglet Design Methods

While Whitcomb generated the framework for the design and implementation of winglets to aircraft in the 1970s [1], much has changed since then in the winglet design process. Of particular interest is the work by Maughmer concerning the realm of sailplane winglet design [2]. The early design philosophy was based on a crossover-point methodology that analyzes the trade-off between the reduction in in-

duced drag and the penalty of the added profile drag of the winglet, as shown in Eq. 1:

$$\Delta D_{\text{induced}} + \Delta D_{\text{profile}} = 0 \tag{1}$$

This relation leads to the crossover speed, V_{CR}, given by Eq. 2, which provides an equation for a velocity that satisfies Eq. 1. It is important to note that this relation is a function of the induced-drag factor, $K(h)$, which is primarily a function of the winglet height, h.

$$V_{CR} = \sqrt{\frac{2W}{\rho b}} \sqrt{\frac{\Delta k(h)}{\pi \Delta h \bar{c} C_{D_{P,WL}}}} \tag{2}$$

When designing the winglet, Maughmer originally set the crossover speed, V_{CR}, slightly higher than the cruise speed to ensure that the added profile drag due to the winglet caused no flight performance penalty in the regular operational range. Ultimately, a modified crossover-point method was developed in Ref. 3, which takes into account other geometric effects than just the winglet height.

Recent work by Takenaka et al. [4] provides a foundation of using a multi-objective evolutionary algorithm to design winglets for a transport aircraft outside of conventional design norms. They follow a very complex winglet-design process, which involves expensive high-fidelity Computational Fluid Dynamics (CFD) and Computational Structural Dynamics (CSD) coupled with an evolutionary algorithm and several data mining techniques. It is important to note in their study that only one flight condition was analyzed, i.e. cruise at $M = 0.8$. A total of two design objectives were pursued by Takenaka et al. [4]: i) minimizing block-fuel used and ii) maximizing the take-off weight.

Multi-Objective Evolutionary Algorithms (MOEA)

In essence, an evolutionary algorithm optimization method is a stochastic search method, which computationally simulates the natural evolutionary process, typically drawing from the Darwinian concept of 'Survival of the Fittest'. Evolutionary algorithms are generally comprised of population members that search the possible solution space. These members are manipulated by a set of operators (mutations, mating, etc.) and evaluated by some type of fitness function. This fitness function determines, which members survive and move on to the next generation of the solution. Theoretically, this process is utilized to search a solution space until an optimum is found. In practice, however, a time constraint typically limits the number of function evaluations (NFE), and thus only an 'acceptable' solution is found. Multi-Objective Evolutionary Algorithms are a feasible optimization approach for the design of winglets of sailplanes that have to satisfy a wide range of design objectives. Multi-objective problems lead to a set of solutions rather than one unique solution. This set of solutions is found through the use of Pareto-optimality theory [5]. Pareto-optimal solutions are those, which when evaluated, cannot be improved without adversely affecting at least one of the objectives. The final optimal selection requires the decision maker to choose from the set *a posteriori* and to weigh the tradeoff between the objectives exemplified by the solution space. Evolutionary algorithm optimization methods have several advantages over traditional deterministic optimization methods, such as calculus/gradient-based, Greedy, or hill-climbing, when the solution space is unknown and/or of non-linear form with multiple peaks and valleys. This is indeed the case for winglets that have a multitude of design parameters including height, taper, twist, along with cant, toe, and sweep angles.

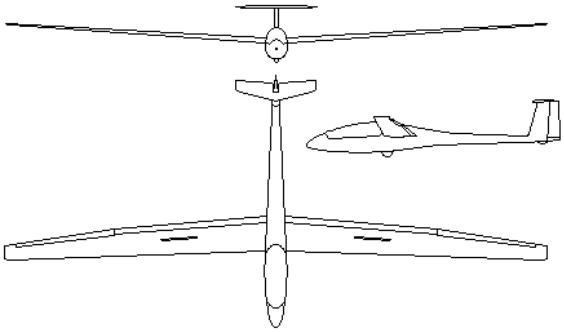

Fig. 1: Sketch of the Janus B sailplane

Contributions of this work

This work presents an efficient computational tool that has the ability to design a winglet that offers performance gains comparable to designs produced using a conventional, i.e. experience-based, design approach. The efficient computational tool that is used in this design process couples an aircraft design code, entitled iFly [6, 7], which uses a higher-order free-wake potential-flow solver entitled FreeWake [8], with the multi-objective evolutionary algorithm optimizer ε-MOEA [9] (epsilon-dominance Multi-Objective Evolutionary Algorithm).

Two primary studies are considered that provide a direct application of the coupled design code/optimizer package to the design of a winglet for the Janus B sailplane, which is a 1978 sailplane design without winglets, see Fig. 1. The first study lays the foundation for implementing a multi-objective optimizer by investigating the feasibility of integrating a single-objective evolutionary algorithm, entitled CMAES (Covariance Matrix Adaptation Evolutionary Strategy) [10] with the aircraft design code iFly [6, 7]. The objective of this initial study is to gain experience with the integrated design approach of flight-performance prediction and evolutionary optimization algorithms. In general, modern sailplanes have to perform well at high angles of attack during thermalling and high-speed flight during cruise between thermals. These diverse requirements lead to the second part of the study that uses a multi-objective optimization algorithm that is based on an ε-MOEA/performance prediction tool.

Using this framework, the second study initially explores two design objectives by optimizing for i) minimizing total aircraft drag at high-speed cruise, and ii) minimizing total aircraft drag during thermalling flight. After this initial study is concluded, a third objective is added as iii) minimizing the root bending moment addition due to a winglet. With these set objectives, the evolutionary algorithm optimizer yields a family of non-dominated solutions, also referred to as a Pareto-optimal set, from which the user performs a selection *a posteriori*. This second study provides insight into the state-of-the-art of using a multi-objective evolutionary algorithm optimizer that is tightly coupled with an aircraft design code. While design objectives ii) and iii) may not be the most relevant objectives regarding winglet designs for competition sailplanes, the focus on objectives other than maximizing the average cross-country speed allows the exploration and demonstration of the multi-objective design optimization. Furthermore, it allows greater flexibility in response to the changing competition rules that allow varying wingtip devices. The value of this work lies in the fact that a suite of advanced winglet designs can be obtained in a reasonable time with the proposed methodology.

Numerical Methods

This work applies a coupled methodology of an aircraft design code and a multi-objective evolutionary algorithm optimizer. The methods are coupled through a simple C script. Significant attention was given to the selection of the individual parts. The selected performance prediction tool is a three-dimensional potential-flow solver entitled FreeWake that was developed by Bramesfeld and Maughmer [8]. The performance code is a subset of an aircraft design tool entitled iFly developed by Kody and Bramesfeld [6, 7]. The selected optimizers are the Covariance Matrix Adaptation (CMA) Evolutionary Strategy by Hansen and Ostermeier [10] and the ε-MOEA (epsilon-dominance multi-objective evolutionary algorithm) by Deb et al. [9]. In the following, each component of the overall coupled method is introduced.

Aircraft Design Code, iFly

Since evolutionary algorithm optimization methods typically require thousands of function evaluations, selecting a performance prediction method that is both time efficient and accurate is an important task. Furthermore, this code is intended for users who may not have access to high-performance computer clusters. The need for fast turnaround times of simulations eliminates the need for high-fidelity CFD simulations for the performance prediction method due to the associated higher computational cost and hardware-infrastructure requirements. Therefore, a three-dimensional potential-flow solver was chosen that had previously been validated and that is capable of resolving the complex flow around a winglet. This is the reason for coupling the selected evolutionary algorithms, which are further described below, with the aircraft design code iFly [6, 7]. A central element to the aircraft performance model is the aerodynamic prediction module rooted in a potential-flow method called FreeWake, developed by Bramesfeld and Maughmer [8]. FreeWake predicts the spanwise lift distribution and induced drag of the wing and empennage. Based on the spanwise lift-distribution prediction and the local chord-based Reynolds numbers along the wing span, the profile drag is determined from look-up tables. A simple stall prediction model, outlined by McCormick [11], was added to the potential-flow code and enhances the accuracy of the aerodynamic solution with insignificant computational cost. In addition, the potential-flow solution finds a trim solution for the horizontal tail, thus effectively capturing any trim-drag effects.

The particular approach in FreeWake uses elements with distributed vorticity to represent the wing and its wake. For the wing model, the spanwise distribution of the bound circulation is described using a second-order spline along each spanwise element. In a sense, the modeling of the bound circulation is very similar to that of Horstmann's Mutliple-Lifting Line method [12], although FreeWake has a trailing-edge condition. In addition, the wake consists of a continuous vortex sheet with a spanwise vorticity distribution that varies linearly along each element. Even as a continuous vortex sheet, the wake can be of non-planar shape and, for example, model the rollup of the shear layer downstream of a wing. For the present study a prescribed model was used that extended the trailing edge along the free stream velocity vector in order to minimize the computational effort. Experience indicates negligible differences in performance predictions for aspect ratios typical for sailplanes. A further strength of the employed potential-flow method lies in its continuous nature that significantly reduces computational costs for load resolutions similar to that of other potential methods, for example vortex-lattice methods. In addition, the continuous model of the wake vorticity shows improved numerical stability due to the absence of singularities as are commonly encountered with conventional vortex-lattice or panel codes.

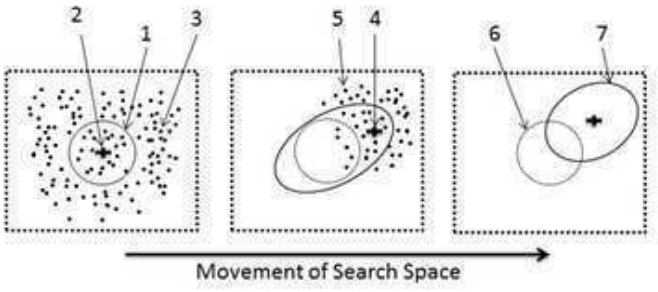

Fig. 2: Concept of the Covariance Matrix Evolutionary Strategy (CMAES). 1) Original Search Radius, 2) Original Centroid, 3) Initial Population Distribution, 4) Moved Centroid, 5) 'Selected' Population Members, 6) Old Search Radius, 7) New Search Radius.

In addition to drag from the wing and empennage, the contributions of other aircraft elements are predicted using either corrected flat-plate skin friction approximations or other semi-empirical approaches. For example, the fuselage drag is estimated using a strip method to determine the skin-friction drag of the exposed surface area. Other correction factors account for pressure-drag and interference losses. The airfoil data used are based on experimental results in Ref. 13 of the FX 67-K-170 airfoil that is used on the Janus B. Various flap settings and Reynolds numbers ranging from 700,000 to 3,000,000 were considered. The winglet airfoil used was the PSU 94-097 airfoil, an airfoil specifically designed for low-speed winglets and tested at The Pennsylvania State University Low-Speed, Low-Turbulence Wind Tunnel at Reynolds numbers ranging from 240,000 to 1,000,000 [14].

Covariance Matrix Adaptation Evolutionary Strategy (CMAES)

The CMA Evolutionary Strategy (CMAES), devised by Hansen and Ostermeier [10], is a stochastic method for real-parameter optimization that relies on an adaptive covariance matrix to overcome many of the hurdles that exist for typical Newton or gradient-based optimizers, for example badly-scaled, non-linear, and highly non-separable objective functions. The overall idea of CMAES is that it utilizes this adaptive matrix to update its search space based on the previous generation's best population members. It also stores information about the movement of the population. With this information, CMAES generates an adaptive path, which better controls the convergence of the solution. The CMAES is also designed to perform very well with small population sizes (≈ 10). Therefore, it is advantageous to problems with large objective function simulation times.

The basics of the CMAES optimizer can be summarized in one simple figure. Figure 2 shows three basic steps for the CMA strategy. The first step, shown in (a), organizes information about the current or initial population of search points. The second step performs a selection and recombination based on the previous generation, as shown in (b). Also shown in (b) is the movement of the mean value of the search distribution (represented by the cross). The third step, shown in (c), exemplifies the actual movement of the search radius. This process essentially repeats and searches the solution space for the global minimum. Unlike the multi-objective function described below, the single-objective CMAES provides only one solution for one objective to the

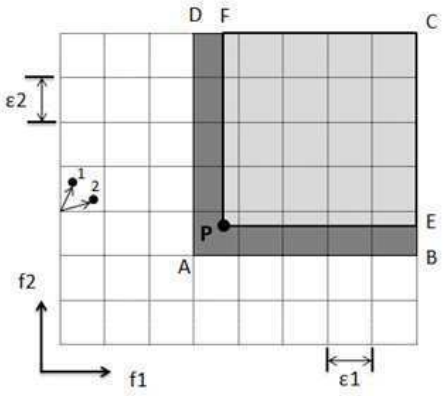

Fig. 3: The ε-Dominance concept

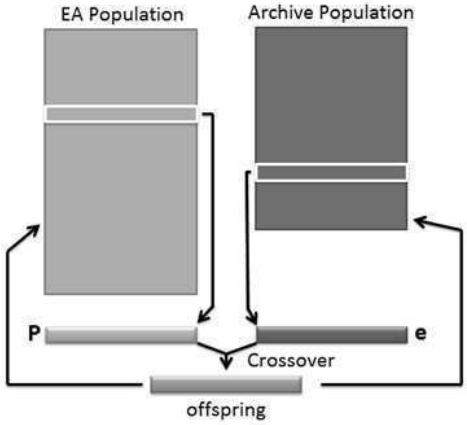

Fig. 4: The ε-MOEA procedure

user. While this removes the decision-making process, which is advantageous for non-experienced designers, it solely optimizes the design for one objective.

Epsilon-Dominance Multi-Objective Evolutionary Algorithm (ε-MOEA)

As stated previously, a multi-objective evolutionary optimizer provides a significant advantage over a single-objective optimizer due to the fact that multi-objective optimizers are able to handle a search and optimization problem with multiple conflicting objectives. A multi-objective optimizer also generates a family of non-dominated solutions, otherwise known as the Pareto-optimal solutions (Pareto-optimal front), which provide the designer with a relationship between the multiple conflicting objectives rather than one single solution. Taking advantage of the Pareto-optimal solutions produced by the multi-objective optimizer allows the designer to avoid a preference-based and weighted system that would prefer one objective over the others prior to the simulation.

Within the field of MOEAs, there are two distinct goals that the algorithm should meet: i) convergence to the true Pareto-optimal front, and ii) generate a well-distributed set of non-dominated solutions. All MOEAs satisfy these requirements, however, most MOEAs are some type of a compromise between providing a well-distributed set of solutions and the speed of convergence for a specific application. Therefore, it would seem very advantageous to find an appropriate MOEA, which provides a good compromise between these two spectrums and the problem at hand. This logic is what led to the decision to using ε-MOEA for the winglet-design problem as computational efficiency is an important factor during a design process and so is ensuring that all design possibilities are fairly analyzed.

The ε-MOEA is a steady-state multi-objective evolutionary algorithm based on the ε-dominance concept that was introduced by Laumanns et al. [15]. This efficient approach provides the user with a well-distributed set of solutions. Within the algorithm, the search space is divided into a number of grids of size (hyper-boxes) that maintain diversity by only allowing one solution to occupy a hyper grid at a given time. The ε-dominance concept is easily explained by Fig. 3. In a minimization problem of two objective functions, f1 and f2, a solution P dominates an entire region PECFP, however, P also ε-dominates the region ABCDA. The ε-dominance concept essentially encompasses all of the ε-sized hyper-boxes that P's traditional domination intersects. Furthermore, if there are two solutions, 1 and 2, within the same hyper-box, the solution closest to the minimization goal is the only survivor.

With an understanding of the ε-dominance concept, the ε-MOEA procedure can be explained. Two co-evolving populations are featured in the ε-MOEA: i) a traditional evolutionary algorithm (EA) population P(t), and ii) an archive population E(t), as seen in Fig. 4. As the simulation begins, the initial EA population, P(0), is created, and the archive population, E(0), is generated by the ε-non-dominated solutions of P(0). Next, two solutions, one from each population, are chosen at random for mating to generate one or multiple offspring. Now the offspring are compared to the EA and archive populations to judge whether or not they should be included to move onto the next generation. In order for the offspring to enter the EA population, the offspring has to dominate, in the traditional sense, one or more of the existing populations. If this is the case, the offspring replaces one of the original population members it dominated. If the offspring, on the other hand, is dominated by any one member of the population, then the offspring is not accepted into the populations. Entering the smaller archive population relies on an ε-dominance comparison in a similar manner as done for the EA population comparison. If the offspring is ε-dominated by the archive population, then the offspring is not allowed to enter, however, if the offspring member ε-dominates the archive member, then the offspring replaces the old archive member. This procedure is simulated for a specified number of iterations or until convergence is achieved. At that point, the archive population members are reported as the final solutions. These final solutions form the ε-non-dominated Pareto-optimal front.

Results

Single-Objective Winglet Design with CMAES

Before taking the full step to solving the multi-objective nature of winglet design, a single-objective optimization was performed of minimizing the total drag of the Janus B sailplane with the addition of a winglet during fast cruise settings. The fast cruise speed was determined through consultation with Janus B owners. Based on these discussions, the selected fast cruise speed was 100 knots (51.4 m/s). At this speed, a typical flap setting of -4° is applied to the FX 67-K-170 airfoil on the main wing. The fast cruise objective was selected based on

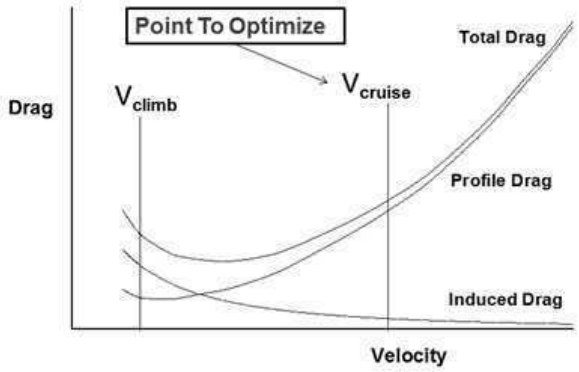

Fig. 5: Single-objective selection on flight envelope

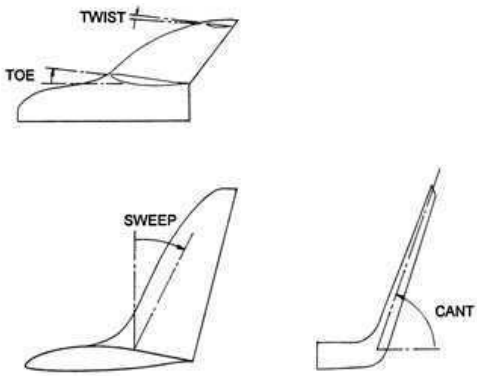

Fig. 6: Winglet angles [2]

Maughmer's design methodology of designing the winglet to not hurt performance at the fast cruise limit, which ensures that the winglet will have a positive influence on the aircraft during its entire flight envelope. The significance with respect to the overall flight envelope and the primary contribution of profile drag for this design point are illustrated in Fig. 5. At lower speeds than the optimization point, induced drag will be larger and an increasingly favorable impact of the winglet can be expected. For the present study, a total of ten parameters were selected to be optimized. Figure 6 illustrates four of these winglet parameters. The remaining parameters, i.e. the winglet height, chord, and taper ratio, are not included in the figure. Figure 7 shows four sections that constitute the geometry specifications for the optimization problem: Section 1 is the panel that serves as an initial transition from the existing main wing

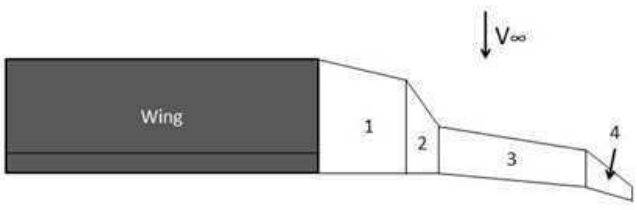

Fig. 7: Winglet geometric sections

Table 1: Janus B simulation configuration

Wingspan	18.20 m
Wing Area	16.94 m^2
Wing Loading	26.8 kg/m^2
Weight	453.61kg/4448 N
Wing Airfoil	FX 67-K-170
Winglet Airfoil	PSU 94-097

Table 2: Variable bounds for optimization design space

Variable	Lower Bound	Upper Bound
Transition span, m	0.061	0.35
Height, m	0.1524	1.524
Cant angle, deg	45	90
Root chord, m	0.1524	0.479
Taper section 3	0.3	1.0
Sweep section 3, deg	0	25
Taper section 4	0.1	1.0
Sweep section 4, deg	0	30
Toe angle, deg	-5	5
Twist, deg	0	10

to the winglet, Section 2 provides the sharpest transition between the winglet and wing and has half of the cant angle than the actual winglet, Section 3 is the first panel of the winglet whose height is limited to 80% of the total winglet height, and Section 4 serves as the winglet tip and comprises 20% of the total height. Table 1 highlights important information regarding the Janus B used in the simulations.

Problem Formulation

In order to achieve the overall goal of finding the optimal winglet geometry, which yields the minimum drag at the aircraft's cruise speed, a parametric study was performed to determine the effects of population size and of the initial random seed values. Of these two experiments, the one that is predicted to have the predominant effect on the solution is population size. Therefore, an analysis was run for populations of 8, 10, 25, 40, and 100. It is important to keep in mind that when reaching a population size of 100, computation time becomes a major factor for the analysis. Therefore, a maximum run-time of 12-hours (which leads to a Number of Function Evaluations, NFE $\approx 20,000$) was set as an upper limit for all simulations. The parametric study also looked at very small population sizes since CMAES is reportedly well conditioned to run at these population sizes. The final experiment involved a random seed analysis of the initial variable dimensions. Four random seed analyses were run to see the effect of variable initial values.

An important factor in this problem formulation is to define the upper and lower bounds for each variable, see Table 2. While the bounded search space is rather large in most cases, this problem's search domain is relatively small due to the real-world constraints of this problem. While several of these bounds were determined by trial and error, such as height, tapers, toe angle (corresponding to the zero lift AoA), and twist, the remaining ones were selected due to aerodynamic considerations or real-world constraints. For example, the transition span's lower bound was set to 0.06 m. (0.2 ft.) for manufacturability reasons, and its upper bound was set due to span limitations. The 0.35 m. (1.148 ft.) was the distance of the wing that was 'chopped' off to make room for

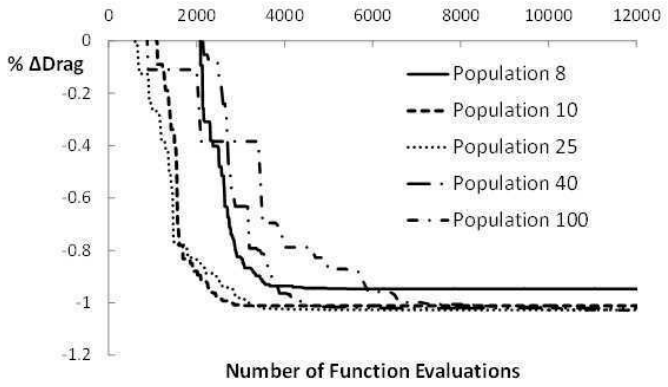

Fig. 8: Single-objective population study

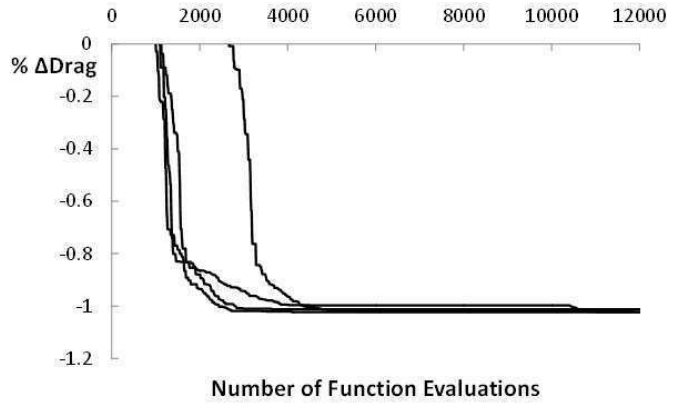

Fig. 9: Single-objective random seed analysis (population size = 10)

the winglet itself. This is standard practice when designing winglets for existing sailplanes due to existing span restrictions. The winglet root chord's lower bound was determined due to manufacturability reasons, and the upper bound was set to the tip chord of the existing wing design. The sweep of sections 3 and 4 were both kept less than 30° in order to prevent cross-flow boundary-layer transition.

CMAES Results: Single-Objective Optimization

The first study focused on the influence of varying population size on the final single-objective solution. The goal of this study was to determine which population produced the best results within the NFE limit, and to determine which population converged fastest to the solution. The objective value used in all simulations was the change in total drag of the aircraft from the conventional configuration (no winglet) to the configuration with the winglet. Therefore, a negative delta drag value corresponds to a reduction in drag. Figure 8 provides insight regarding the NFE that each population requires to reach the optimal solution. Figure 8 also shows that the lower populations achieve a lower drag value in fewer NFE than the larger population runs, although it should be noted that the smallest population of 8 misses the optimum solution by less than a tenth of a percentage point. Furthermore, the change in drag seems to level before 7000 NFE for all population cases.

The second experiment looked at the effect of a changing random seed, which defines the initial starting point of the solution. Figure 9 shows the effect of four random seed solutions that were run. This experiment was conducted to study the effect of how a changing random seed value would affect convergence and if different solutions were found (various local minima, etc.). As can be seen, this random seed study showed very little variation in the minimal value attained, while only slight differences in convergence speed are seen. This holds true even in the case where due to the chosen starting point it takes over 2500 NFE to indicate any drag savings. This study was run with a population size of 10.

Figure 9 shows that a drag reduction of 1.03% (Drag = -3.13 N) is achieved during the cruise-flight condition (100 kt, -4° flap). This means that the winglet operates successfully at the rest of the flight envelope, even if it is not optimized for those other flight conditions. Drag reduction due to the winglet during high-lift coefficient thermalling flight (50 kt, +8° flap, 45° bank), which will be introduced later, resulted in a drag reduction of 1.41% (Drag = -2.42 N). This is a very small change in drag, which is expected due to the fact that designing a winglet for

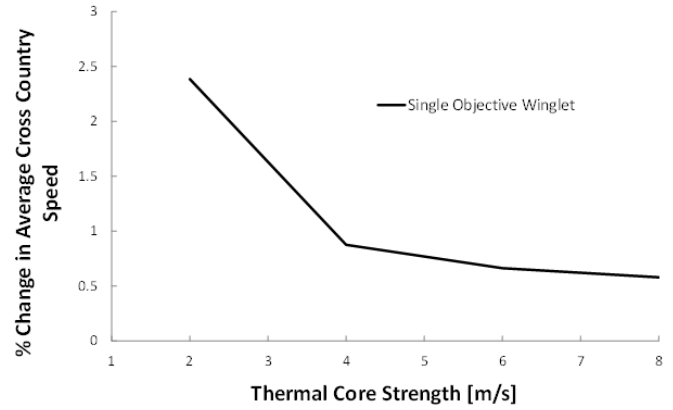

Fig. 10: Percent change in cross-country speed between baseline Janus B and single-objective optimized winglet versus the thermal core strength.

cruise only demands minimal area, therefore minimal profile drag. Interestingly, the focus on cruise also appears to favour a cant angle of nearly 48 degrees. The cant angle is a compromise of the need to reduce the profile-drag penalty of the vertical winglet, while securing whatever small savings of the relatively small induced drag savings of the nonplanar geometry. Figure 10 provides information about how the optimized winglet effects the average cross-country speed of the glider. The average cross-country speed is increased 2.4% and 0.6% for a weak (2 m/s core strength) and strong (8 m/s core strength) thermals, respectively. Table 3 provides the final variable values found in this minimization study, which used a population of 10. Additionally, Fig. 11 provides a visual creation of the final winglet. It is important to recall that the toe angle is defined with respect to the aerodynamic AoA. The PSU 094-97 airfoil has a zero-lift angle of attack that hovers around -4.5° with respect to the airfoil chord for most Reynolds numbers. Therefore, a -1.8° aerodynamic angle of attack toe angle actually refers to a +2.7° geometric toe out.

Multi-Objective Winglet Design with ε-MOEA

With the preliminary single-objective study complete, the foundation was laid to move on to a multi-objective study. As has been reiterated

Table 3: Final winglet dimensions from single objective study

Transition span, m	0.061
Height, m	0.416
Cant angle, deg	48.5
Root chord, m	0.1524
Taper section 3	0.35
Sweep section 3, deg	0.22
Taper section 4	0.96
Sweep section 4, deg	16.69
Toe angle, deg	-1.8
Twist, deg	1.26

Fig. 11: Final winglet obtained from single-objective study

before, the winglet design problem for a sailplane is multi-objective in nature. Therefore, utilizing the unique abilities of a MOEA seems rather appropriate. Before this study commenced, proper attention was given to ensure the problem formulation was laid out correctly. A two-objective study was thoroughly investigated: i) minimize total drag of the aircraft at fast cruise, ii) minimize total drag of the aircraft at a standard thermalling velocity. A fast cruise speed of 100 kt (51.4 m/s) was selected as in the single-objective study. This also corresponded to the same -4° flap setting of the Janus B's airfoil. The thermal velocity was selected to be 50 kt (25.7 m/s) with a load factor of 1.41 (bank angle of 45°). A positive 8° flap setting is typically used by pilots at this flight condition. This velocity was determined by discussing common thermalling speeds with Janus B owners. Figure 12 shows the flight envelope with the two objectives highlighted.

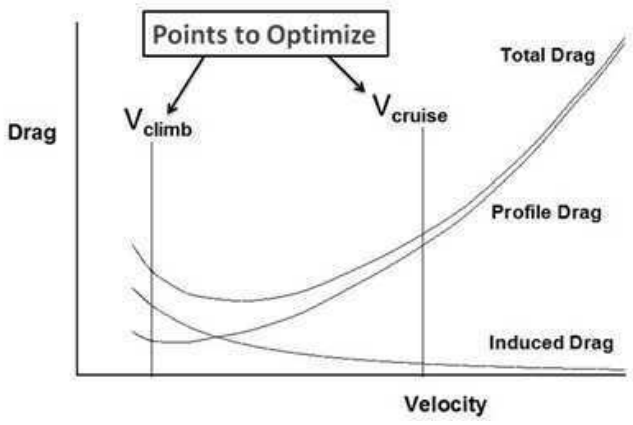

Fig. 12: Multi-objective selection on flight envelope

With these two objectives, particular attention was given to the population size, the selected random seed, the crossover probability, the crossover distribution, and the bounds of the variables. It was expected that, since two objectives were being addressed in this study, the required simulation time would have to be significantly increased. Therefore, the maximum simulation time allowed was increased to 24 hours.

All simulations were run on a single core of a 2.8 GHz quad core processor running Linux. At the end of this two-objective study, a third objective was added that focused on the minimization of the root bending moment addition due to the winglet. This was added due to the fact that since the Janus B has an existing structure that was originally designed to support the design loads without a winglet.

Problem Formulation

The geometric construction of the optimization problem from the single-objective study was also used for the multi-objective studies. This was done in order to gain insights into the effect that each objective has on the final solutions. As introduced earlier, four experiments were performed to determine acceptable optimizer settings for this specific problem. At first, a population study was performed by looking at populations of 50, 75, 100, and 200. Unlike CMAES, the ε-MOEA does not work efficiently with small populations. The four population sizes considered were based on standards recommended by the developer of the optimizer. A typical random seed analysis was also performed to ensure that it does not have a major impact on the final Pareto front solution. A crossover distribution study was also performed with values of 10 and 15. Only two values were considered because they were the two recommended settings. Finally, a crossover-probability study was performed at two different population values. The investigated crossover values were 0.6, 0.8, 0.9, and 1.0.

Three constraints were used on this problem's formulation. Most importantly, a span constraint was employed. Originally, the Janus B sailplane had a span of 18.2 m. Therefore, no population member was able to produce results, if it had a span greater or less than the original 18.2 m. The second constraint was on the tip chord of the winglet. This was set due to manufacturability considerations and allowed only designs with a tip chord greater than 5.08 cm to be included in the population. The third constraint enforced the common-sense rule that the sweep of the winglet tip section should be greater than the sweep of the first winglet section.

ε-MOEA Results: Two-Objective Optimization

In the following sections, the optimal values are given in the form of Pareto fronts. For the present two-objective study, the y-axis objective is the percent change of the aircraft drag during fast cruise, and the x-axis represents the percent change in drag during a thermalling flight condition. In order for a winglet to successfully 'work', it should produce a reduction (negative change) in drag for both objectives. Only a portion of the non-dominated population members that comprise the Pareto front will successfully 'work' according to this definition. It is important to remember that each of these population members on the front represents a unique winglet design for the Janus B. The final selection process will be addressed later, as it is highly dependent on the designers' goals and preferences.

In order to generate the proper settings for this type of problem using the ε-MOEA, several studies were performed prior to the final analysis. Before the optimizer is run, a random seed value, between 0 and 1, is selected to seed the initial population used in the algorithm. Five random seed values were investigated for a population of 100, and the

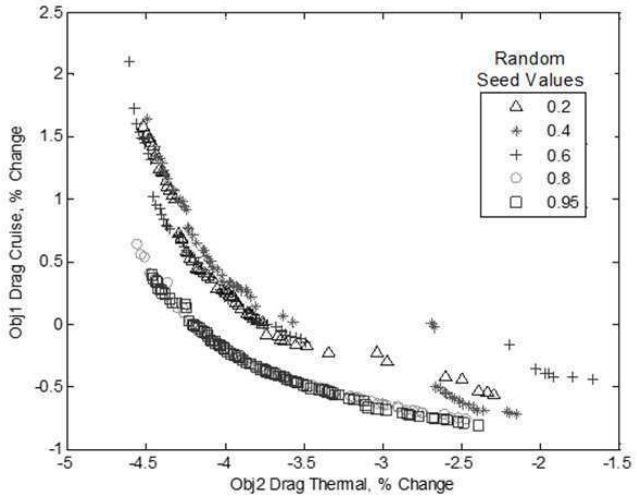

Fig. 13: Multi-objective random seed analysis. (population size = 100)

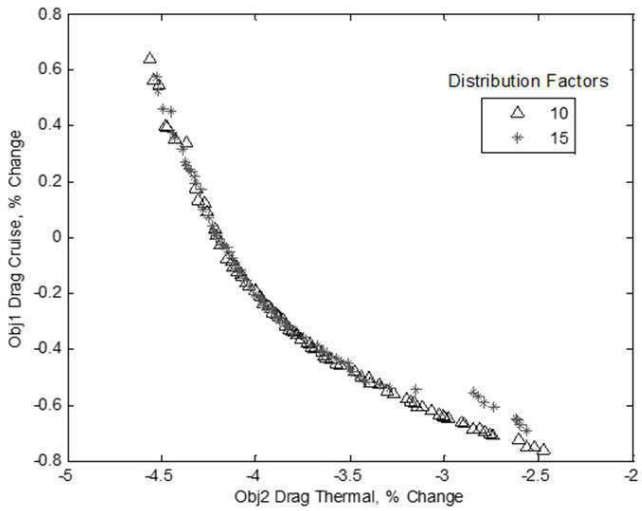

Fig. 15: Multi-objective crossover distribution analysis. (population size = 100)

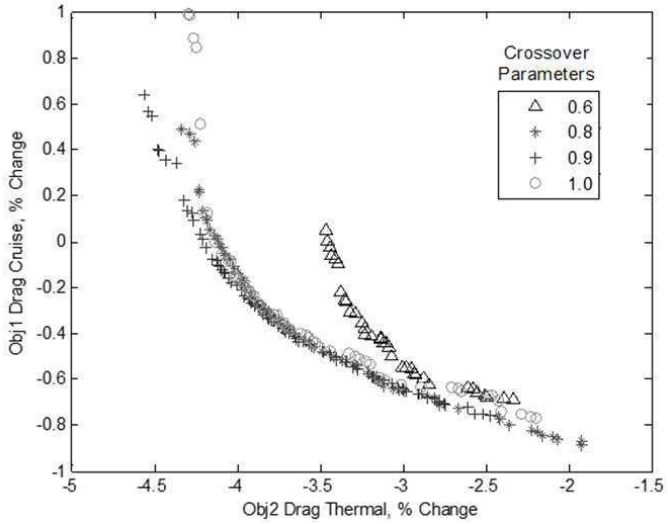

Fig. 14: Multi-objective crossover probability analysis. (population size = 100)

simulations were allowed to last for 10,000 NFE (23.5 hours). All other settings were kept as defaults for this analysis. Of the five seed values, it can be seen in Fig. 13 that higher random seed values provide more 'optimal' Pareto fronts. Ultimately, a value of 0.8 was selected for the rest of the analysis, since it was the originally recommended value and because it performed among the top of the tested values.

Next, a suitable setting for the crossover probability was determined. This setting describes how often a population member is forced to mate with another population member during the analysis. Available settings range between 0.6 and 1.0. Four values were investigated within these constraints, see Fig. 14. In the end, a crossover probability of 0.9 proved to generate the best Pareto front for the given constraints. It can be rationalized that the reason why a higher crossover probability produces better results at the 10,000 NFE mark is due to the fact that the population is being forced to search the solution space at a more rapid pace. The cases of lower crossover probability are more prone to keep older population members in the analysis for more generations, which slows the search speed of the optimizer. The population and function-evaluation limit were kept the same as for the random seed study. A crossover probability test was not run for the CMAES optimizer due to the fact that mutation and mating is controlled in a different way than for the ε-MOEA solver.

Crossover distribution refers to how a population member has a mating partner selected for itself. As for the crossover distribution, solver guidelines recommend either a value of 10 or 15. As can be seen in Fig. 15, both 10 and 15 result in relatively similar Pareto fronts for these settings. However, the value of 10 shows a slightly more optimal front and was therefore selected.

Finally, and most importantly, a population study was performed. Selecting the correct population setting is always a compromise between a well-spread solution and the performance of the Pareto front after 10,000 NFE. The Pareto fronts seen in Fig. 16 highlight this trade-off. The Pareto front associated with a population of 200 clearly results in the Pareto front that contains the most wide-spread solutions. However, it is also one of the worst in terms of convergence relative to the Pareto fronts obtained for other population sizes. A population of 100 was selected because its solution provided the best compromise between a well-spread front and relative convergence for a given NFE limit.

The final two-objective Pareto front is shown in Fig. 17. As was expected, the final Pareto front produces a well-spread solution that contains population members, i.e. winglet designs, which work at both conditions as well as at only one condition. As can be expected, a design that reduces drag at the cruise objective does the same for the thermalling objective. However, this does not hold true the other way around. Almost half of the solutions on the Pareto front yield increases in drag during cruise but significant decreases in drag during thermalling. A more detailed investigation of how individual variables change over this front may provide insight into why some population members work efficiently and others seemingly do not. Highlighting the effect of each variable would be beyond the scope of this paper. Therefore, only two variables will be discussed, i.e. height and root chord of the winglet, as these two variables are the strongest drivers in

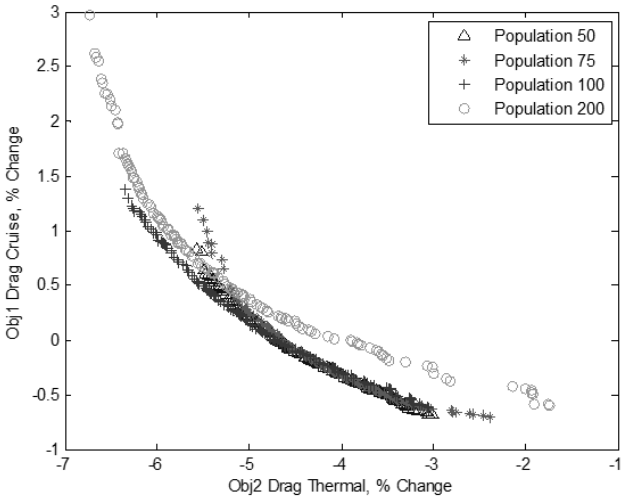

Fig. 16: Multi-objective population study

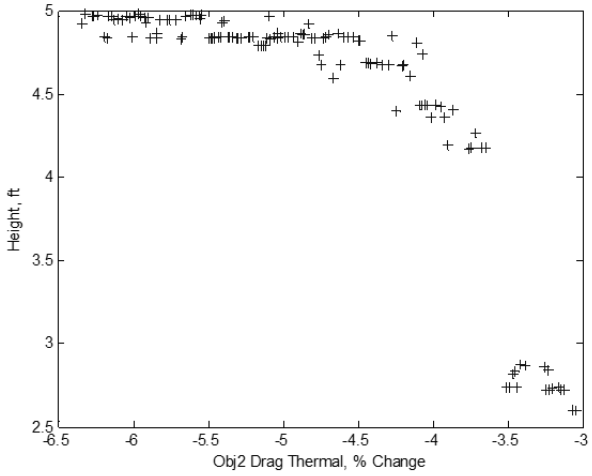

Fig. 18: Effect of winglet height on two-objective Pareto front

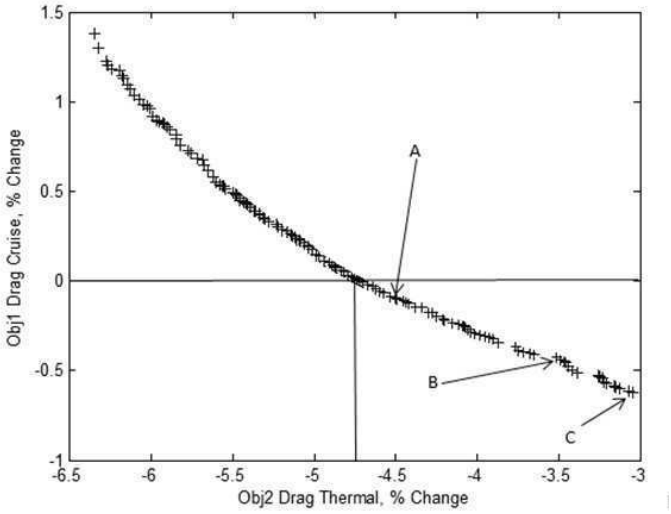

Fig. 17: Final two-objective Pareto front including selections of winglets A, B, and C

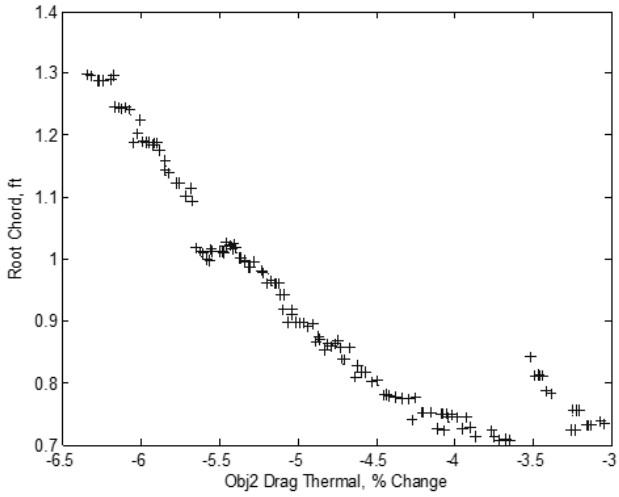

Fig. 19: Effect of winglet root chord on two-objective Pareto front

affecting the final winglet performance. The results shown in Figs. 18 and 19 help shed light on the importance of these variables. The figures show the individual make-up of each population member by setting the *x*-axis to the percent change of the thermal objective (Objective 2), and by setting the *y*-axis to the variable value associated with the Pareto-front population member at that value of the thermal objective. In these figures, it can be seen that the variation in height and chord lead to significant changes in thermal drag savings, more so than all of the other variables. From these figures, two simple conclusions can be drawn: i) in order to achieve a larger reduction in drag during thermalling, a taller winglet and a larger root chord are necessary, ii) to produce a more effective winglet at high speeds, a smaller winglet with a smaller root chord (less wetted area) is advantageous. The latter observation is also exemplified by the single-objective study. While the root chord is gradually changing over the varying thermal objective, the height wants to remain as large as possible until thermal objective values of -3.5% are reached. This is a very interesting result and is most likely attributed to the fact that an increase in bending moment is not taken into account. Since this winglet is being designed for an existing sailplane, it is beneficial to minimize the addition to the root bending moment due to the winglet for structural considerations. Since this objective is not yet in place, the optimizer is unable to determine that winglet heights of 1.52m are most likely not feasible and simply impractical. Before the bending-moment objective is added, an example design selection from the Pareto front is performed.

An *a posteriori* selection is highly dependent on the designer's preferences. If the designer values drag reductions at cruise speed over thermal speeds, it is apparent how to select a design with those attributes. The same holds for a designer looking for a large increase in thermal performance. In order to exemplify these extremes, three different designs are selected and sketched to provide guidance for future design

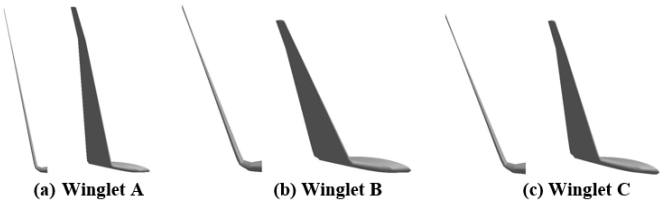

Fig. 20: Selected winglets obtained from two-objective study

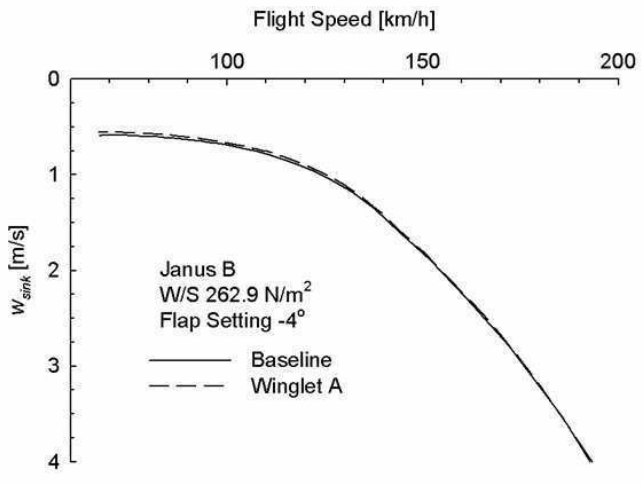

Fig. 21: Speed polar of baseline and configuration with winglet A

selections. Figure 17 has three population members selected, i.e. A, B, and C designs. Designs A and C are meant to cover the extremes of the 'working' winglet possibilities, while design B provides a compromise between the two. The sketches of these winglet design possibilities are illustrated in Fig. 20. The speed polars of the baseline and winglet A are shown in Fig. 21. Small differences are visible in the low speed region, which is in line with the data presented in Fig. 17. Although the speed polars of winglets B and C are virtually indistinguishable from winglet A, the differences in the subsequent cross-country speeds are pronounced in Fig. 22. All three winglet designs provided the large performance advantages over the baseline during weak thermalling conditions with thermal core strengths of less than 4 m/s. Although less pronounced, they are all able to maintain that advantage even during stronger thermals. Please note that thermal core strength is the vertical velocity of the air, which can differ considerably from the actual climb rate. For example, the configuration with winglet A achieves climb rates of 0.91 m/s and 6.17 m/s for thermal-core strengths of 2 m/s and 8 m/s, respectively. Winglet A, whose dimensions are listed in Table 4 together with those of winglets B and C, is considered to have the best performance improvement with respect to cross-country speed. This suggests that more emphasis should be given to minimizing drag during thermalling conditions while still ensuring that no increase in drag occurs at high cruise speeds.

ε-MOEA Results: Three-Objective Optimization

With the addition of the bending-moment objective, further insight is expected regarding the large winglet heights seen on the final two-objective Pareto front. The resulting Pareto front from the three-objective study using a population size of 100 can be seen in Figs. 23 and 24 after 10,000 NFE. Since it is difficult to realize a three-dimensional plot on paper, two projected side views are used in the following. The traditional Obj1. vs. Obj2. plot is shown in Fig. 23, while the effect of the root bending moment on the Pareto surface is highlighted in Fig. 24. As expected, winglets that yield improved thermal performance produce larger root bending moments. When comparing cruise and thermal drag tradeoffs, however, configurations with thermal drag reductions of about 2% to 4% seem to also have benefits in cruise drag. As was shown for the two-objective study, Fig. 25 illustrates how the winglet height is affected by a change in the Pareto front for the three-objective case. Overall, the winglet heights are decreased significantly to around 3.5 feet (1.07m) compared to the two-objective case with nearly 5 feet (1.55m). Based on these criteria three winglets, D, E, and F, were chosen as examples. Their thermal and cruise drag changes over the baseline are indicated in Fig. 23. Their geometric data are summarized in Table 5. The speed polar of winglet D is shown in comparison to the baseline in Fig. 26. Just as with the two-objective results, the speed polars of the other two designs are very similar, yet the average cross-country speed differs, especially when modeling weak weather conditions as shown in Fig. 27. Overall, winglet D seems to provide the best cross-country performance under weak weather conditions and

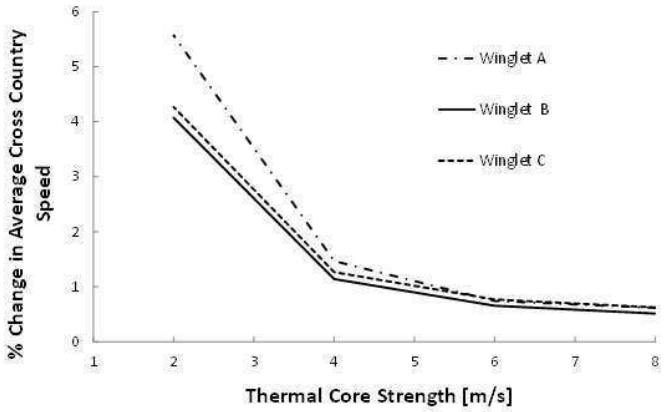

Fig. 22: Percent change in cross-country speed between baseline Janus B and selected winglets A, B, and C versus the thermal core strength

Table 4: Final dimensions of winglet A, B, and C from two-objective study

Variable	Winglet A	Winglet B	Winglet C
Transition span, m	0.0609	0.0611	0.0620
Height, m	1.469	0.835	0.793
Cant angle, deg	78.3	69.1	67.9
Root chord, m	0.245	0.257	0.223
Taper section 3	0.305	0.333	0.312
Sweep section 3, deg	1.69	6.79	2.92
Taper section 4	0.698	0.688	0.767
Sweep section 4, deg	7.96	8.56	7.67
Toe angle, deg	-3.41	-2.13	-2.12
Twist, deg	1.12	1.18	1.90

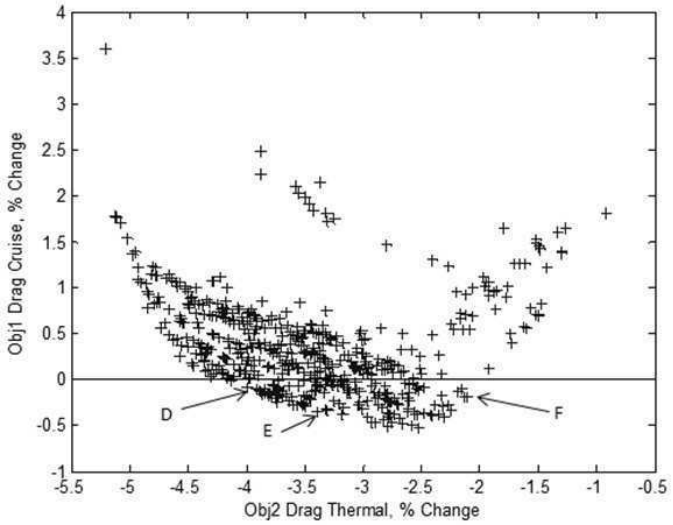

Fig. 23: Obj. 1 vs. Obj. 2 of the final three-objective Pareto front including selections of winglets E, D, and F

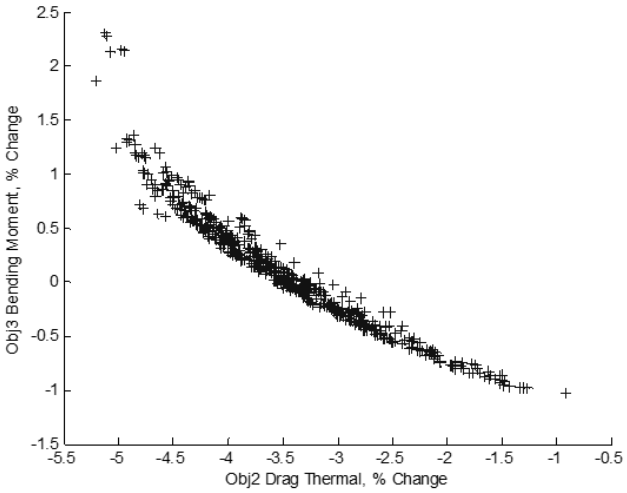

Fig. 24: Obj. 3 vs. Obj. 2 of the final three-objective Pareto front

Table 5: Final dimensions of winglet D, E, and F from three-objective study

Variable	Winglet A	Winglet B	Winglet C
Transition span, m	0.0609	0.0615	0.0709
Height, m	1.087	1.042	1.001
Cant angle, deg	74.1	74.6	74.5
Root chord, m	0.227	0.200	0.154
Taper section 3	0.48	0.51	0.46
Sweep section 3, deg	2.6	1.3	5.2
Taper section 4	0.54	0.55	0.82
Sweep section 4, deg	6.3	9.5	9.2
Toe angle, deg	-1.60	-1.70	1.06
Twist, deg	4.89	3.80	4.54

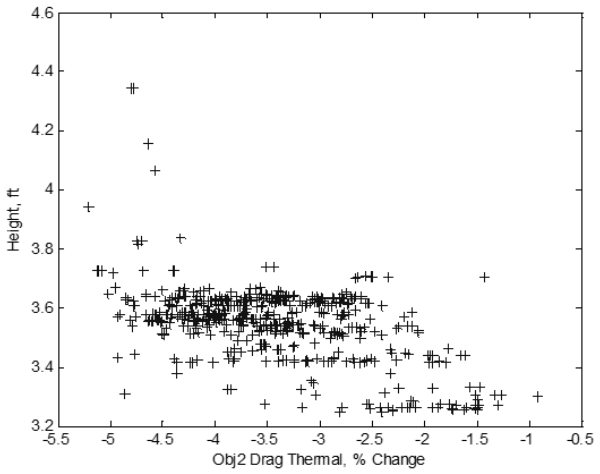

Fig. 25: Effect of winglet height on three-objective Pareto front

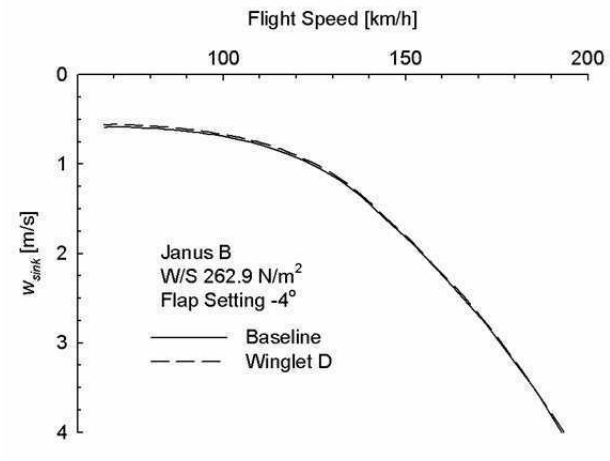

Fig. 26: Speed polars of baseline and configuration with winglet D

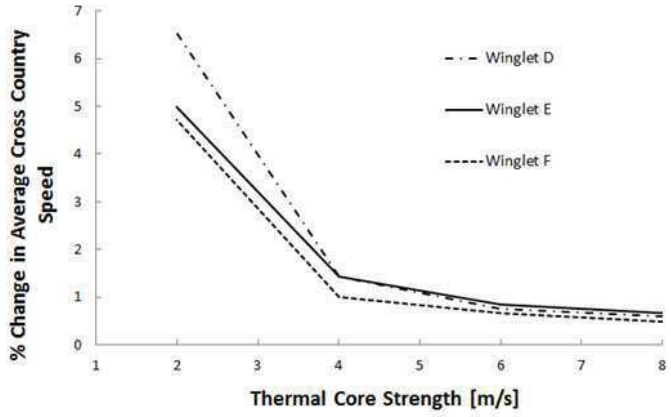

Fig. 27: Speed polars of baseline and configuration with winglet D

without any significant penalties when stronger thermals are modeled, although interestingly winglet E seems to perform slightly better with thermal core strengths in excess of 4 m/s. Nevertheless, even though it is about 40 cm shorter than winglet A, winglet D exhibits a faster average cross-country speed over the entire modeled thermal-core strength range. At this point of the study, no distinct characteristic can be associated with this performance advantage. The outcome of this study, however, is a reflection of the complexity of the possible design space that is only very limitedly described with the chosen three objectives.

Conclusions

An aircraft design code, iFly, that uses an efficient and accurate potential-flow solver based on distributed vorticity elements was coupled with the single-objective CMAES and multi-objective ε-MOEA evolutionary strategies with the intent to study the feasibility of using such a coupled methodology for advanced winglet design. A total of three objectives were considered: i) minimizing total aircraft drag during high-speed cruise, ii) minimizing total aircraft drag during high-lift thermalling flight, and iii) minimizing the root bending moment addition due to the winglet. While this combination of objectives may not be the most relevant and pertinent for competition sailplane design, the main goal of this work was to explore and demonstrate that evolutionary-based design and optimization of sailplane winglets is feasible on today's desktop computers.

Differences between the final winglet designs from the single-to multi-objective studies are apparent. Each time an objective was added, the resulting Pareto front reflected that change. Moving from the single-objective study to the multi-objective study, it becomes evident that the average cross-country speed is affected more beneficially by reducing the thermalling drag while maintaining near constant cruise drag compared to solely focusing on reducing cruise drag. Introducing the third objective of wing-bending moment leads to limitations in winglet height without necessarily resulting in performance limitations.

The importance of thermal drag with respect to average cross-country speed compared to cruise-drag changes is also supported by the findings of the multi-objective study. For example winglet A, which primarily reduces thermalling drag with negligible influence on cruise performance, has significantly higher gains in average cross-country speed during weak weather conditions in comparison with winglets B and C, which lead to smaller gains during thermalling, but larger ones during cruise. These differences, however, diminish with increasing thermal-core strengths, although all three designs still have noticeable advantages over the baseline.

Overall, results obtained are promising with respect to the implementation of evolutionary algorithm optimizers for sailplane winglet design. It is important to note that these designs were found during simulations that took less than 24 hours to complete on a conventional desktop computer. Future work will be directed towards the performance of the computed designs compared to existing winglets as well as the addition of other direct objectives such as the average cross-country speed.

Acknowledgment

The authors gratefully acknowledge the Department of Aerospace Engineering at the Pennsylvania State University for providing computing resources used in this work.

References

[1] Whitcomb, R. T., "A Design Approach and Selected Wind-Tunnel Results at High Subsonic Speeds for Wing-Tip Mounted Winglets," Tech. Rep. TN D-8260, NASA, July 1976.

[2] Maughmer, M. D., "Design of Winglets for High-Performance Sailplanes," *Journal of Aircraft*, Vol. 40, No. 6, 2003, pp. 1099–1106.

[3] Maughmer, M. D. and Kunz, P., "Sailplane Winglet Design," *Proceedings of the XXV OSTIV Congress*, OSTIV, St. Auban, France, 1997.

[4] Takenaka, K., Hatanaka, K., Yamazaki, W., and Nakahashi, K., "Multidisciplinary Design Exploration of a Winglet," *Journal of Aircraft*, Vol. 45, No. 5, 2008, pp. 1601–1611.

[5] Coello, C., Lamont, G., and Van Veldhuizen, D., *Evolutionary Algorithms for Solving Multi-Objective Problems*, Science and Business Media, Springer, 2nd ed., 2007.

[6] Kody, F., "An Integrated Aircraft Design and Performance Prediction Tool — Design, Validation, and Demonstration," *Aerospace Sciences Meeting*, AIAA, Nashville, Tennessee, 2012, Paper No. 2012-0126.

[7] Kody, F. and Bramesfeld, G., "Small UAV Design Using an Integrated Design Tool," *International Journal of Micro Air Vehicles*, Vol. 4, No. 2, 2012.

[8] Bramesfeld, G. and Maughmer, M. D., "Relaxed-Wake Vortex-Lattice Method Using Distributed Vorticity Elements," *Journal of Aircraft*, Vol. 45, No. 2, 2008, pp. 560–568.

[9] Deb, K., Mohan, M., and Mishra, S., "A Fast Multi-Objective Evolutionary Algorithm for Finding Well-Spread Pareto-Optimal Solutions," Tech. Rep. 2003002, Indian Institute of Technology, 2003, KanGAL.

[10] Hansen, N. and Ostermeier, A., "Adapting Arbitrary Normal Mutation Distributions in Evolution Strategies: The Covariance Matrix Adaptation," *International Conference on Evolutionary Computation*, IEEE, 1996, pp. 312–317.

[11] McCormick, B. W., *Aerodynamics, Aeronautics and Flight Mechanics*, John Wiley & Sons, Hoboken, New Jersey, 1995.

[12] Horstmann, K. H., *Ein Mehrfach-Traglinienverfahren und seine Verwendung für Entwurf und Nachrechnung nichtplanarer Flügelanordnungen*, Ph.D. thesis, DFVLR, Institut für Entwurfsaerodynamik, Braunschweig, Germany, 1987.

[13] Althaus, D., *Stuttgarter Profilkatalog*, Vieweg, 1981.

[14] Maughmer, M. D., Swan, T., and Willits, S., "The Design and Testing of a Winglet Airfoil for Low Speed Aircraft," *19th Applied Aerodynamics Conference*, AIAA, Anaheim, California, 2001.

[15] Laumanns, M., Thiele, L., Deb, K., and Zitzler, E., "Combining Convergence and Diversity in Evolutionary Multi-Objective Optimization," *Evolutionary Computing*, Vol. 10, No. 3, 2002, pp. 263–282.