

# Optimization of Non-Conventional Airfoils for Martian Rotorcraft using Direct Numerical Simulations

L. Caros\*, O. Buxton\* and P. Vincent\*

Corresponding author: lidia.caros-roca19@imperial.ac.uk

\* Imperial College London, London SW7 2AZ, UK.

**Abstract:** Designing helicopters to fly on Mars is a challenging problem. The Martian atmosphere is characterized by a very low density and lower speed of sound compared to Earth. These conditions require the rotor blades of Martian rotorcraft to operate in a low-Reynolds-number (1,000 to 10,000) compressible flow regime, not typical of conventional helicopters. Several studies have shown a performance drop when conventional airfoils are used in these conditions. Non-conventional airfoils with sharp leading edges and flat surfaces show an improved performance compared to conventional airfoils. In order to optimize such unconventional airfoils, several studies have made use of evolutionary techniques. These techniques usually require many cost function evaluations, and hence, they typically employ Reynolds-Averaged Navier Stokes (RANS) solvers because of their low computational cost. However, RANS solvers have limited predictive capability when the flow becomes unsteady and separated at moderate angles of attack, in particular at low Reynolds numbers. The current study overcomes this limitation by performing the optimization with high-order accurate direct numerical simulations (DNS) using the compressible flow solver in PyFR ([www.pyfr.org](http://www.pyfr.org)). Specifically, we optimize triangular airfoil shapes with two-dimensional (2D) and three-dimensional (3D) DNS at two different angles of attack  $\alpha$ . At  $\alpha = 6^\circ$ , results from optimization using 2D DNS and 3D DNS are similar because the flow is predominantly two-dimensional. At  $\alpha = 12^\circ$ , results from optimization with 2D DNS and 3D DNS are different due to the three-dimensional breakdown of coherent vortices in the span-wise direction, which can not be captured by 2D simulations.

*Keywords:* Martian Aerodynamics, Optimization, Direct Numerical Simulations, Low-Reynolds-Number, Unsteady Flows.

## 1 Introduction

On April 2021, Ingenuity performed the first-ever heavier-than-air flight on Mars [1] and has since then achieved 28 successful flights in the Martian atmosphere. Ingenuity is an autonomous rotorcraft with two counter-rotating rotors with conventional airfoil profiles [2]. The rotorcraft has had to overcome the many challenges of flying in the Martian atmosphere. These include Mars' very thin atmosphere, with a density lower than 1% of that on Earth, and Mars' surface temperature and atmospheric composition, which cause a lower speed of sound than on Earth, limiting the rotor speeds.

For over a 100 years, designers have worked to optimize airfoils for the full range of terrestrial conditions. However, flow conditions on Mars involve a low Reynolds number ( $\approx 10^3$  to  $10^4$ ) compressible regime, at which conventional airfoils lose performance. The design of airfoils for Martian atmospheric conditions is a new and emerging field, and optimal airfoil shapes are still to be developed. Studies by Munday et al. [3] and Konig et al. [4] have achieved promising results using airfoils with sharp leading edges and flat surfaces. Sharp-leading-edge airfoils induce a higher adverse pressure gradient than conventional airfoils, causing an undulation of the separated shear layer that leads to a roll-up of coherent vortices on the suction side of the airfoil, which improves the airfoil's performance under these conditions.

In order to find the optimal sharp-leading-edge airfoil shapes, several studies have used evolutionary algorithms such as Genetic Algorithms (GA) [5, 6]. These algorithms usually involve many cost function evaluations, and hence they typically employ Reynolds-Averaged Navier Stokes (RANS) solvers because of their low computational cost. However, RANS solvers have limited predictive capability when the flow becomes unsteady and separated at moderate angles of attack, thus limiting the utility of the approach.

Recent advances in solver technology and GPU hardware have enabled the use of high-order accurate direct numerical simulations (DNS) for airfoil shape optimization under Martian conditions. Previous work by Caros et al. [7] has shown the utility of using PyFR [8] to perform DNS for non-conventional airfoils, obtaining accurate results. Additionally, Caros et al. [7] show the effect of incorporating more representative flow physics in high-fidelity simulations of the flow over an unconventional airfoil at different angles of attack. Results showed the need to resolve the three-dimensional flow structures at moderate to high angles of attack when the flow separates from the sharp leading edge.

In this work, we undertake optimization using high-fidelity DNS via the compressible flow solver in PyFR [8]. Specifically, we optimize a triangular airfoil under the conditions of chord-based Reynolds number  $Re = 3,000$  and Mach number  $M = 0.15$  with two-dimensional (2D) and three-dimensional (3D) DNS at two different angles of attack,  $\alpha = 6^\circ$  and  $\alpha = 12^\circ$ .

## 2 Methodology

### 2.1 Computational Fluid Dynamics Setup

In the present work, we use two- and three-dimensional high-order accurate Direct Numerical Simulations (DNS) with PyFR to evaluate the aerodynamic performance of different unconventional airfoil shapes.

The simulations ran on Nvidia Tesla P100-16GB GPUs on Piz Daint at the Swiss National Supercomputing Centre (CSCS).

#### 2.1.1 Solver and Numerical Method

PyFR is a cross-platform solver based on the high-order Flux Reconstruction approach of Huynh [9]. In the present work, fourth-order polynomials were used to represent the solution within each element of the mesh, thus nominally achieving fifth-order accuracy in space. A Rusanov Riemann solver was used to calculate the inter-element inviscid fluxes, the Local Discontinuous Galerkin approach was used to calculate the inter-element viscous fluxes, an explicit RK45 scheme [10] was employed to advance the solution in time, no anti-aliasing was employed, and all runs were performed using double-precision arithmetic.

#### 2.1.2 Computational domain

The airfoil to be optimized is a triangular airfoil with a chord of 1.0 and infinitely sharp edges. The computational mesh is adapted autonomously to each airfoil shape produced by the optimizer. The mesh used for the particular triangular airfoil that has been validated in Caros et al. [7] is shown in Figure 1. The mesh is comprised of both structured and unstructured regions. A structured quadrilateral mesh is located adjacent to the airfoil surfaces, and the discretization within this region is adapted to the geometry changes by maintaining the cell size below the one used for the validated airfoil. An additional structured quadrilateral mesh is located in the wake region, extending a distance of 6.0 downstream of the airfoil. The remainder of the streamwise-vertical plane is tessellated with an unstructured mix of quadrilaterals and triangles. The 2D meshes have approximately  $\sim 6 \times 10^3$  elements and  $\sim 16 \times 10^3$  degrees of freedom per equation in total. For the 3D simulations, the two-dimensional mesh is extruded a length of 0.6 and discretized in the span-wise  $z$  direction, generating hexahedra and triangular prisms throughout the domain. The 3D meshes have  $\sim 155 \times 10^3$  elements and  $\sim 19.4 \times 10^6$  degrees of freedom per equation in total.

#### 2.1.3 Boundary conditions

A characteristic boundary condition is applied at the  $x = -10$  inflow,  $x = 20$  outflow,  $y = -10$  bottom, and  $y = 10$  top boundaries, with a prescribed density  $\rho = 1.0$ , velocity  $\mathbf{v} = \{v_x, v_y\} = \{1.0, 0.0\}$  and pressure




Figure 1: Computational mesh used for the two-dimensional DNS of a triangular airfoil with  $x_a = 0.3$  and  $y_a = 0.05$ . Reused with permission from L. Caros. *Copyright 2022, Lidia Caros.*

$p = 31.746$  to achieve  $M = 0.15$  and  $Re = 3,000$ . At the airfoil surface, an impermeable no-slip adiabatic boundary condition is prescribed. No turbulence is injected at the inflow.

#### 2.1.4 Data extraction

At  $t = 0$ , the simulations are initiated with initial conditions of uniform density  $\rho = 1.0$ , velocity  $\mathbf{v} = \{v_x, v_y\} = \{1.0, 0.0\}$  and pressure  $p = 31.746$  throughout the domain. Simulations are advanced a period  $t_t$  to remove initial transients.  $t_t = 50t_c$  for the 2D simulations when  $\alpha = 6^\circ$  and  $\alpha = 12^\circ$ , with  $t_c = c/u_\infty$  where  $u_\infty$  is the time-averaged  $x$ -velocity magnitude measured at the centre of the inflow boundary. For the 3D simulations,  $t_t = 60t_c$  when  $\alpha = 6^\circ$  and  $t_t = 90t_c$  when  $\alpha = 12^\circ$ . For the first  $5t_c$  of each simulation the velocity prescribed at the inflow plane is modified as follows  $\mathbf{v} = \{1 + 0.2 \sin(100t), 0.2 \sin(100t)\}$  in order to trigger flow instabilities. Simulations are then advanced a further  $t_a$  during which data is extracted for analysis. For both  $\alpha = 6^\circ$  and  $\alpha = 12^\circ$ ,  $t_a = 50t_c$  for the 2D simulations and  $t_a = 60t_c$  for the 3D simulations.

The quantities extracted from the simulations are the lift and drag coefficients, defined as

$$C_L = \frac{F_L}{q_\infty c}, \quad C_D = \frac{F_D}{q_\infty c},$$

respectively, where  $F_L$  and  $F_D$  are the time-averaged lift and drag forces, respectively,  $c$  is the airfoil's chord, and  $q_\infty$  is the time-averaged dynamic pressure.  $F_L$  and  $F_D$  are obtained by time-averaging the sum of viscous and pressure forces on the airfoil in the  $y$  and  $x$  directions, respectively, over  $t_a$  and  $q_\infty$  is obtained via

$$q_\infty = \frac{1}{2} \rho_\infty v_\infty^2$$

where  $\rho_\infty$  is the density measured at the centre of the inflow plane and time-averaged over  $t_a$ .

## 2.2 Optimization Setup

Genetic Algorithms (GA) [11] are used in this work to optimize triangular airfoil shapes. These algorithms are based on the theory of evolution and make use of genetic operators to optimize towards the global optimum. GAs explore the design space to find the fittest set of design variables, for the given objective functions, by evolving populations of individuals through generations. Although these algorithms require multiple evaluations to converge, they are considered suitable for problems with multiple objectives and nonlinear objective spaces with various local optima. The particular GA used is the Non-dominated Sorting

Genetic Algorithm II (NSGA-II) by Deb et al. [12]. The optimization targets minimizing two objective functions and thus, outputs a set of non-dominated solutions, a Pareto front.

The optimization process consists of a set of executable scripts that autonomously interact with each other. The main optimization runs on Intel Xeon E5-2650 v3 2.30GHz CPUs on Piz Daint at the Swiss National Supercomputing Centre (CSCS).

### 2.2.1 Numerical method

Pymoo [13], an open-source Python-based framework for single- and multi-objective optimization, is used for the optimization. The framework allows customization, which facilitates its coupling with other frameworks, such as PyFR and Gmsh.

Gmsh is used for generating the automated mesh for each of the airfoil shapes proposed by the optimizer. Gmsh is an open-source meshing tool that can be used to create unstructured curved element meshes of complex geometries in a semi-automated fashion. A notable feature of Gmsh is that it can be scripted via a Python API.

### 2.2.2 Design Space and Objective Space

The design variables are the  $x$  and  $y$  coordinates of the apex of the triangular airfoil,  $x_a$  and  $y_a$ , respectively, as shown in Figure 2. The coordinates are measured from the airfoil's leading edge with an angle of attack of  $\alpha = 0^\circ$ . The design space is constrained to avoid unrealistic geometries, such as unfeasible thicknesses, and issues with the mesh, such as skewed elements.




Figure 2: Schematic of the triangular airfoil with design variables  $x_a$  and  $y_a$ .

The objective functions are maximization of lift and minimization of drag, evaluated by their non-dimensional coefficients,  $C_L$  and  $C_D$ , respectively. The bi-objective optimization problem with two constraints is defined as

$$\text{maximize } f_1(\mathbf{x}) = C_L(x_a, y_a), \quad \text{minimize } f_2(\mathbf{x}) = C_D(x_a, y_a)$$

### 2.2.3 Optimization parameters

The optimization is initiated with a population of individuals; each individual is a set of apex coordinates representing an airfoil shape. The first population is seeded with Latin Hypercube Sampling (LHS) of  $x_a$  and  $y_a$  coordinates. The population consists of 30 individuals, with 15 offspring bred at each generation. These numbers were selected based on the parallel computational resources. Genetic operators such as selection, crossover and mutation are applied when producing the offspring to balance exploitation and exploration of the solution space. The optimization finishes when the termination criterion is met; in this case, it is set as a specific number of generations. The optimization convergence is evaluated with the hypervolume metric [14], which measures the area dominated by the Pareto front at a particular generation with respect to a reference point. A more detailed explanation of the optimization parameters used in this work can be found in Caros et al. [15].

## 3 Flow physics of a triangular airfoil

At the flow conditions at which Martian helicopters operate, low chord-based Reynolds number (1,000 to 10,000) compressible regime, conventional airfoils experience a performance drop. The drop occurs when the



flow is in a sub-critical regime ( $Re < 50,000$ ), and there is laminar separation without turbulent reattachment. Unconventional airfoils with sharp leading edges outperform conventional airfoils due to the strong pressure gradients that the sharp leading edge induces, which generate vortices that roll up on the suction side of the airfoil as shown in Figure 3. These vortices are formed from the undulation of the separated shear layer caused by the Kelvin-Helmholtz instability.




Figure 3: Instantaneous density gradient magnitude of the flow over a triangular airfoil at an angle of attack of  $\alpha = 6^\circ$ .

Previous work by Munday et al. [3] and Caros et al. [7] on a particular triangular airfoil with the apex at  $x_a = 0.3$  and  $y_a = 0.05$  shows that, under conditions of  $Re = 3,000$  and  $M = 0.15$ , the airfoil has a highly nonlinear lift curve. As the angle of attack increases, various flow regimes can be determined. These include a laminar two-dimensional regime when  $\alpha < 7^\circ$  and a highly unsteady three-dimensional regime when  $\alpha > 9^\circ$ . The changes in the flow regime coincide with the switching of the separation point. For this particular airfoil, at  $\alpha < 7^\circ$ , the flow fully separates from the apex of the triangular airfoil, while at  $\alpha > 9^\circ$ , the flow separates from the leading edge. With the flow separating from the sharp leading edge, larger and stronger vortices roll up on the suction surface compared to those shed from the apex, inducing a large area of low pressure on the suction surface, enhancing lift.

Figure 4a shows plots of lift coefficient for different angles of attack obtained from experiments by Munday et al. [3] and two- and three-dimensional DNS by Caros et al. [7]. The 2D DNS PyFR results are in reasonable agreement with the experimental  $C_L$  for low  $\alpha$ . However, as  $\alpha$  increases and the flow becomes highly separated and unsteady, the 2D PyFR simulations start to significantly over-predict  $C_L$  relative to the experimental results. Results from 3D DNS with periodic boundary conditions in the span-wise direction are similar to 2D DNS when  $\alpha \leq 9^\circ$ . However, at high angles of attack 3D DNS achieves better agreement with experimental data because the 3D span-wise DNS captures the span-wise breakdown of the large vortical structures that are shed above the surface of the airfoil. Caros et al. [7] found that a span of 0.6 is needed to avoid artificially suppressing the spanwise extent of the three-dimensional structures.

The switching of the separation point can be clearly seen in Figures 4b and 4c, which show the time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for  $\alpha = 7^\circ$  in (b) and  $\alpha = 9^\circ$  in (c). When  $\alpha = 7^\circ$ , a small recirculation bubble is formed from the leading edge, but the flow reattaches rapidly, and complete separation occurs from the apex. On the other hand, when  $\alpha = 9^\circ$ , full separation occurs from the leading edge without reattachment.

## 4 Optimization results

The triangular airfoil is optimized under the flow conditions of Mach number  $M = 0.15$  and chord-based Reynolds number  $Re = 3,000$ . Four different optimizations are run, two at each of the angles of attack,  $\alpha = 6^\circ$  and  $\alpha = 12^\circ$ , with two different cost function evaluations each, 2D and 3D DNS with PyFR.

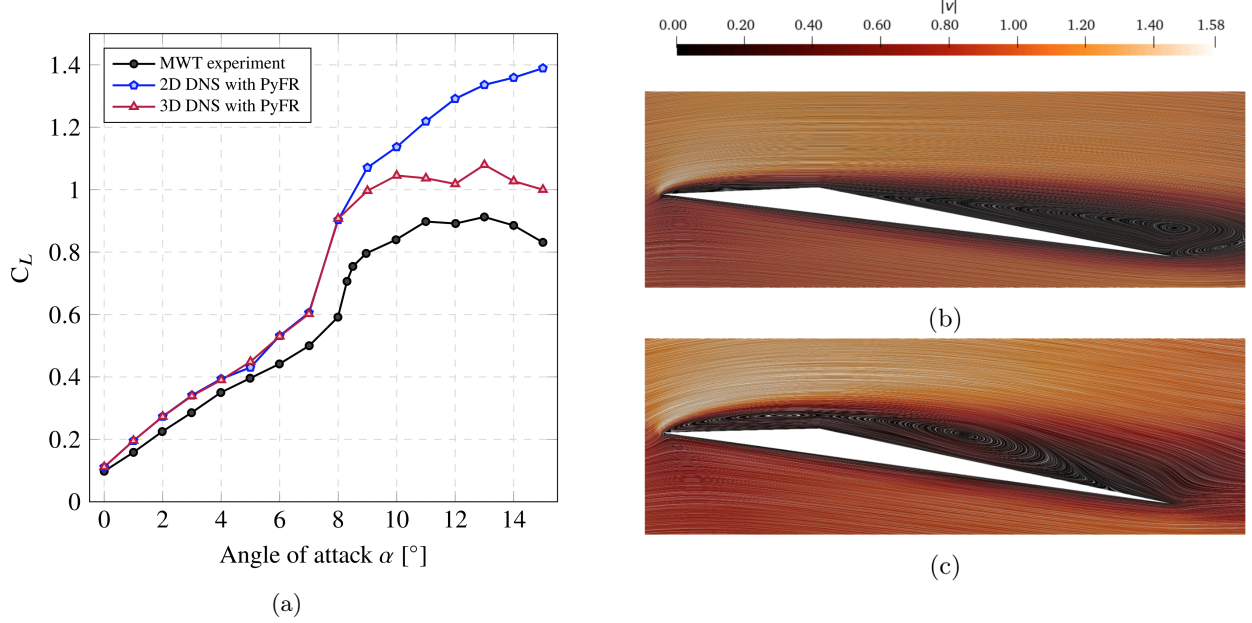


Figure 4: Plots of lift coefficient  $C_L$  as a function of angle of attack  $\alpha$  obtained from the MWT experiments of Munday et al. [3] and 2D and 3D spanwise periodic DNS with PyFR by Caros et al. [7] (a). Images of time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for angles of attack  $\alpha = 7^\circ$  (b) and  $\alpha = 9^\circ$  (c) from 3D spanwise periodic DNS. Reused with permission from L. Caros. Copyright 2022, Lidia Caros.

#### 4.1 Low angle of attack ( $\alpha = 6^\circ$ )

The 2D optimization ran for 15 generations, with 15 individuals in each generation except for the first one, which has 30 individuals, resulting in 240 2D DNS evaluations. Each of the 2D simulations takes approximately 2 hours to run with 1 GPU, hence, the 2D DNS optimization consumes approximately 480 GPUhs. The 3D optimization ran for 10 generations only due to its higher cost, resulting in 165 3D DNS evaluations in total. Each of the 3D simulations takes approximately 24 hours to run with 20 GPUs, therefore, the 3D DNS optimization consumes approximately 80,000 GPUhs.

The convergence of the optimizations was evaluated with the hypervolume [14] performance indicator, metric commonly used when the converged Pareto front is unknown. It calculates the dominated area by the Pareto front at each generation with respect to a reference point. Figure 5 shows that the 2D optimization is converged by generation 13 as the hypervolume only changes slightly with more generations. The 3D DNS optimization can not be considered converged by generation 10, and more generations should be run to assess its convergence.

Both 2D and 3D DNS optimizations start from the same initial population. The first population is seeded with a Latin Hypercube Sampling (LHS) set of airfoil shapes in the constrained design space, as shown in Figure 6a. The aerodynamic coefficients obtained by each airfoil evaluated with 2D and 3D DNS are shown in the objective space in Figure 6b. The objective space shows that there are airfoils that obtain very similar lift and drag coefficients with 2D and 3D simulations, while others differ significantly, mainly in the lift coefficient. Note that this correlates with previous work by Caros et al. [7] but at high angles of attack, not at  $\alpha = 6^\circ$ . Nevertheless, the set of airfoils that are non-dominated, and hence, that form the Pareto front at generation one, do not show significant differences in the aerodynamic coefficients with the two different evaluations.

As the objective functions drive the optimization, and there are some differences in the results obtained from 2D and 3D DNS evaluations, the two optimizations evolve differently by testing different airfoil shapes. However, both optimizations evolve to a very similar Pareto front of non-dominated solutions. Figure 7a shows the optimum airfoils found by both 2D and 3D DNS optimizations at generation 10, and Figure 7b




Figure 5: Hypervolume metric as a function of the number of generations showing the convergence of both 2D and 3D DNS optimizations at  $\alpha = 6^\circ$ .

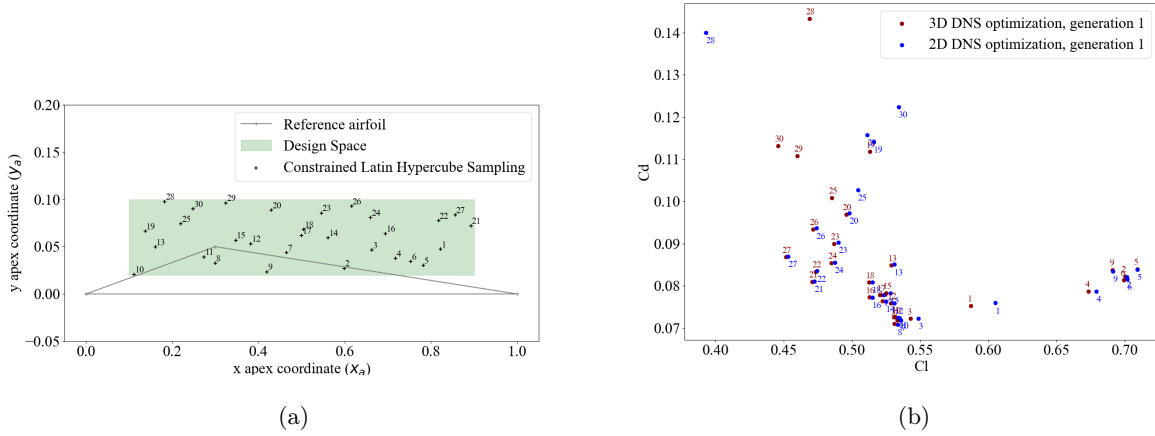


Figure 6: Design space showing the initial sampling of LHS apex coordinates for the 2D and 3D DNS optimizations for  $\alpha = 6^\circ$  (a). Objective space showing the aerodynamic forces obtained from evaluating the initial population of airfoils for  $\alpha = 6^\circ$  with 2D and 3D DNS (b).

shows their corresponding aerodynamic coefficients. In Figure 7a the airfoils with minimum drag (individual 1), maximum lift-to-drag ratio (individuals 13 and 15) and maximum lift (individual 30) are drawn with lines, while the rest are just defined with points in the apex coordinates. Figure 7a shows that there are differences in the optimum apex coordinates found by the 2D and 3D DNS optimizations. This could be caused by the different physics in the two different optimizations or by the stochasticity associated with GAs. Figure 7b also shows, in green, the optimum airfoils found by the 2D DNS optimization evaluated with 3D DNS. It can be seen that these are still optimum as they are within the Pareto front of the 3D DNS optimization. Hence, performing 2D DNS optimization is suitable for a triangular airfoil at an angle of attack of  $\alpha = 6^\circ$ , allowing a reduction in cost of approximately 150 times.

Figure 8 shows the time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for three optimum airfoil shapes found by the 3D DNS shape optimization. Figure 9 shows images of instantaneous Q-criterion,  $Q = \frac{1}{2}(\|\Omega\|^2 - \|S\|^2)$ , isosurfaces of  $Q = 1$ , coloured by velocity magnitude over the same three optimum airfoil shapes. The optimum shapes correspond to the airfoils with minimum drag (a), maximum lift-to-drag ratio (b), and maximum lift (c). Figure 8 (a), which corresponds to the minimum drag airfoil, shows a small laminar separation bubble near the leading edge, and flow separation from the apex of the airfoil. For the maximum lift and lift-to-drag ratio, the optimizer

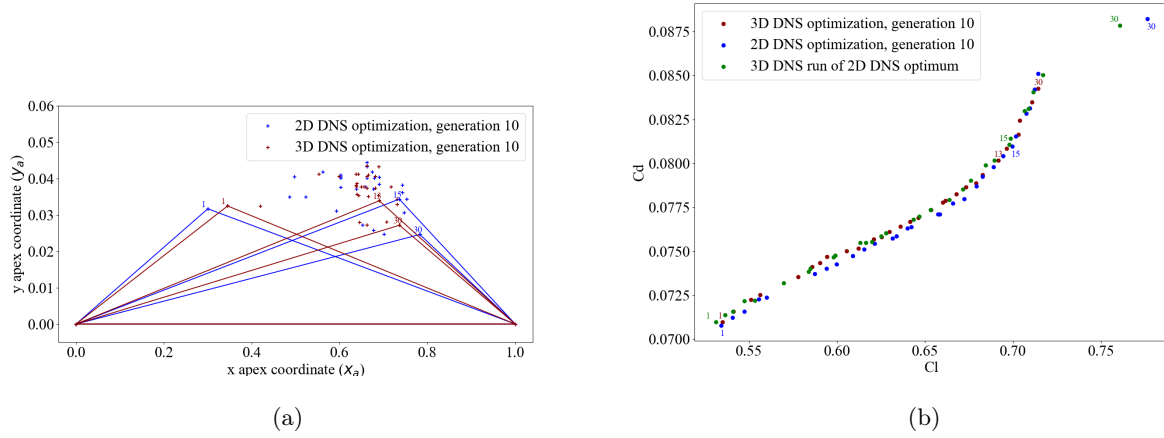


Figure 7: Design space showing the optimum apex coordinates found by the 2D and 3D DNS optimizations at generation 10 for  $\alpha = 6^\circ$  (a). The lines show the geometries of airfoils found with minimum  $C_D$  (1), maximum  $C_L/C_D$  (13 and 15) and maximum  $C_L$  (30). Objective space showing the aerodynamic forces obtained from the optimum airfoils at generation 10 for  $\alpha = 6^\circ$  with 2D and 3D DNS, and 3D DNS runs of the optimum 2D DNS airfoils (b).




Figure 8: Images of time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for triangular airfoils with minimum  $C_D$  (a), maximum  $C_L/C_D$  (b), and maximum  $C_L$  (c) obtained from 3D DNS optimization at generation 10 for  $\alpha = 6^\circ$ .

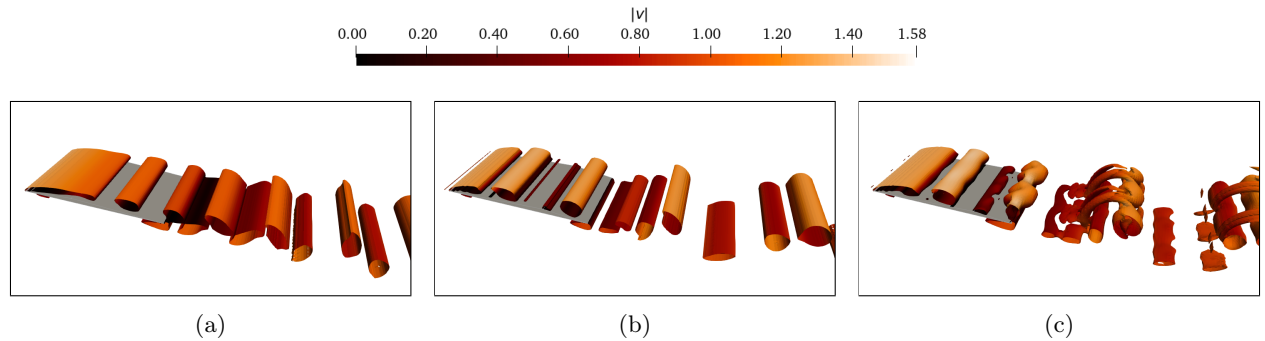


Figure 9: Images of instantaneous Q-criterion isosurfaces  $Q = 1$ , coloured by velocity magnitude  $|\mathbf{v}|$ , for triangular airfoils with minimum  $C_D$  (a), maximum  $C_L/C_D$  (b), and maximum  $C_L$  (c) obtained from 3D DNS optimization at generation 10 for  $\alpha = 6^\circ$ .

has found an apex position that induces the previously described roll-up of vortices from the leading edge. The Q-criterion images in Figure 9 show that, although there is a vortex roll-up from the leading edge for

(b) and (c), the flow remains two-dimensional over most of the airfoil surface with a spanwise domain of 0.6. Three-dimensional structures can be identified towards the trailing edge in Figure 9 (c) for the maximum lift airfoil, which explains the differences in lift coefficient for the maximum lift airfoil found by the 2D DNS optimization.

Given that the three examined airfoils inhabit the extremes in the objective space, it may be presumed that all the airfoils in the Pareto front have similar flow structures, with two-dimensional flow over most of the airfoil surface. This corroborates the ability of using 2D DNS evaluations for triangular airfoil optimization at an angle of attack of  $\alpha = 6^\circ$  and, likely, at lower angles of attack.

## 4.2 High angle of attack ( $\alpha = 12^\circ$ )

The optimizations at angle of attack of  $\alpha = 12^\circ$  ran for more generations as they experienced a more challenging convergence. The 2D optimization ran for 25 generations, with 15 individuals in each generation except for the first population, which has 30 individuals, totaling 390 2D DNS evaluations. Each of the 2D simulations takes approximately 2 hours to run with 1 GPU, hence, the 2D DNS optimization consumes approximately 780 GPUhs. The 3D optimization ran for 15 generations, resulting in 240 3D DNS evaluations. Each of the 3D simulations takes approximately 24 hours to run with 36 GPUs, therefore, the 3D DNS optimization consumes approximately 200,000 GPUhs.

Figure 10 shows the hypervolume metric for both 2D and 3D DNS optimizations for  $\alpha = 12^\circ$  to assess their convergence. The 2D optimization shows convergence from generation 20, while the 3D optimization can not be considered converged, and further generations should be run.

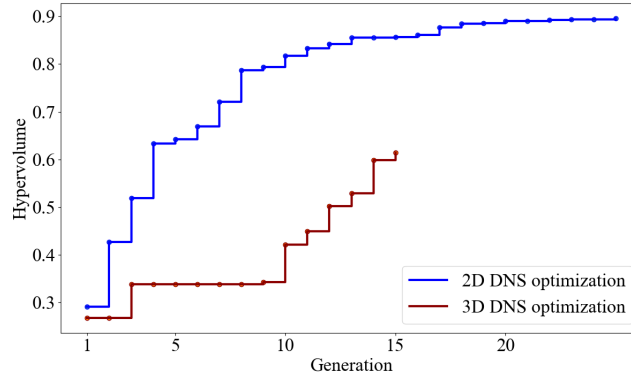


Figure 10: Hypervolume metric as a function of the number of generations showing the convergence of both 2D and 3D DNS optimizations at  $\alpha = 12^\circ$ .

The 2D and 3D DNS optimizations start from the same initial population seeded with a Latin Hypercube Sampling (LHS) set of airfoil shapes in the constrained design space, as shown in Figure 11a. Note that the constrained design space for the current optimizations has been extended compared to  $\alpha = 6^\circ$  since preliminary tests showed airfoils evolving towards the borders of a smaller design space. The aerodynamic coefficients obtained by each of the airfoils with 2D and 3D DNS evaluations are shown in the objective space in Figure 11b. As expected, the aerodynamic forces obtained from the first population of airfoils at an angle of attack of  $\alpha = 12^\circ$  are very different when evaluated with 2D and 3D DNS. The results at  $\alpha = 12^\circ$  from 3D DNS compared to 2D DNS show lower lift and drag coefficients, which agrees with previous work by Caros et al. [7] for a particular triangular airfoil at  $\alpha = 12^\circ$ .

Figure 12b shows the optimum airfoils found by the 2D DNS optimization at generation 20 and 3D DNS optimization at generation 15. Figure 12a shows the corresponding aerodynamic coefficients obtained from evaluating the optimum airfoils, which form the Pareto front. Similar to before, Figure 12b shows three optimum airfoils drawn with lines while the rest are represented with their apex coordinates. The three drawn optimum shapes correspond to the airfoils with minimum drag (individuals 13 and 2), maximum lift (individuals 3 and 24), and maximum lift-to-drag ratio (individual 18) for the 2D DNS optimization. The

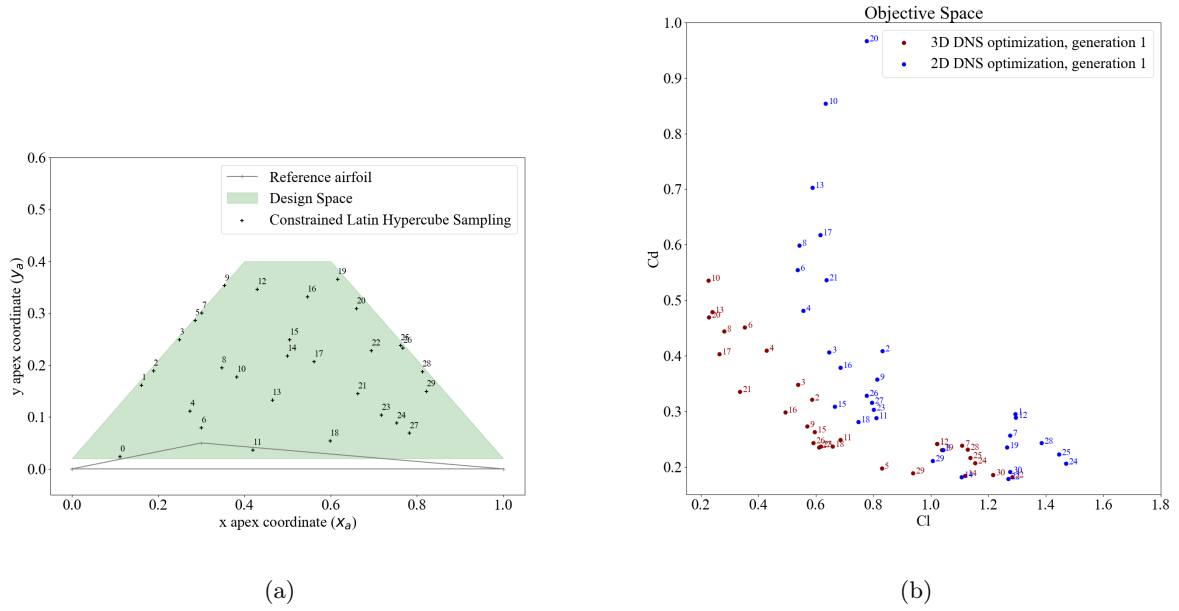


Figure 11: Design space showing the initial sampling of LHS apex coordinates for the 2D and 3D DNS optimizations for  $\alpha = 12^\circ$  (a). Objective space showing the aerodynamic forces obtained from evaluating the initial population of airfoils for  $\alpha = 6^\circ$  with 2D and 3D DNS (b).

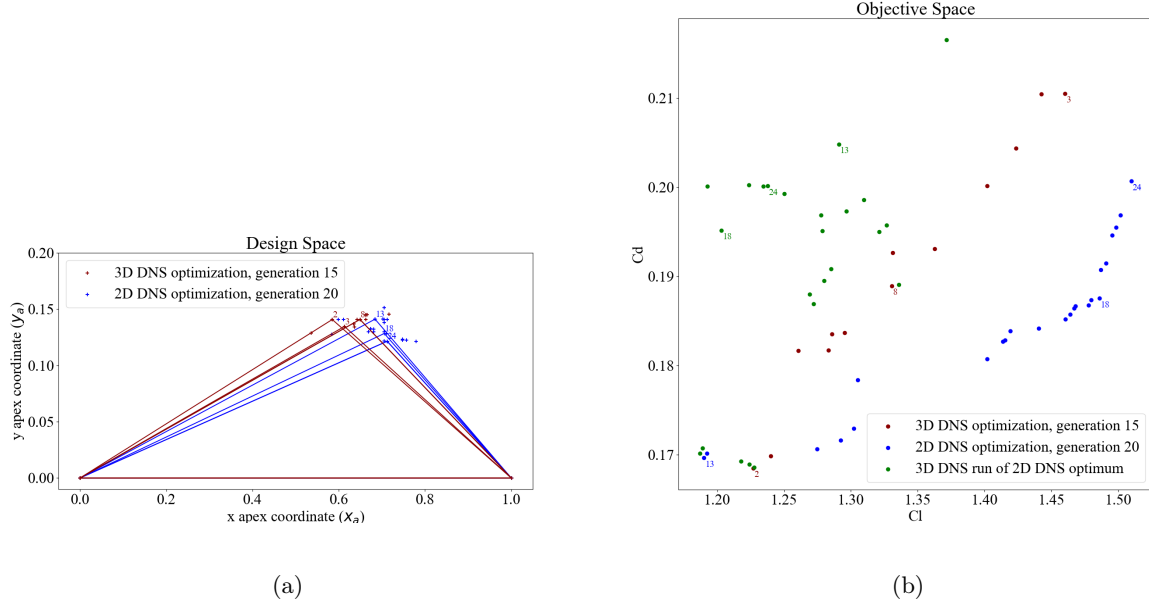


Figure 12: Design space showing the optimum apex coordinates found by the 2D and 3D optimizations for  $\alpha = 12^\circ$ . The lines show the geometries of airfoils found with maximum  $C_L$  (8 and 30), minimum  $C_D$  (1) and maximum  $C_L/C_D$  (5 and 16) (a). Objective space showing the aerodynamic forces obtained from the optimum airfoils for  $\alpha = 12^\circ$  with 2D and 3D DNS, and 3D DNS runs of the optimum 2D DNS airfoils (b).



maximum lift-to-drag ratio airfoil for the 3D DNS optimization coincides with the maximum lift airfoil, so an additional optimum airfoil is plotted (individual 8).

The main visible difference in the optimum airfoil shapes in Figure 12b is that the airfoils obtained from the 2D DNS optimization have their apex coordinates downstream of those obtained from 3D DNS optimization. The objective space in Figure 12a shows that, when optimizing at  $\alpha = 12^\circ$ , the Pareto fronts obtained when using 2D DNS or 3D DNS as cost function evaluation are completely different, with the 2D DNS results giving much higher lift and lower drag. If the airfoils obtained from the 2D DNS optimization are evaluated with 3D DNS, their aerodynamic coefficients are worsened by 5 to 15%, showing that evaluating with 2D causes an over-prediction of the lift and an under-prediction of the drag. The most important outcome from evaluating the 2D DNS optimum with 3D DNS is that the airfoils fall into the dominated area of the 3D DNS Pareto front. Hence, optimizing triangular airfoils at  $\alpha = 12^\circ$  with 3D DNS evaluations is essential.




Figure 13: Images of time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for triangular airfoils with minimum  $C_D$  (a), high  $C_L/C_D$  (b), and maximum  $C_L$  and  $C_L/C_D$  (c) obtained from 3D DNS optimization at generation 15 for  $\alpha = 12^\circ$ .

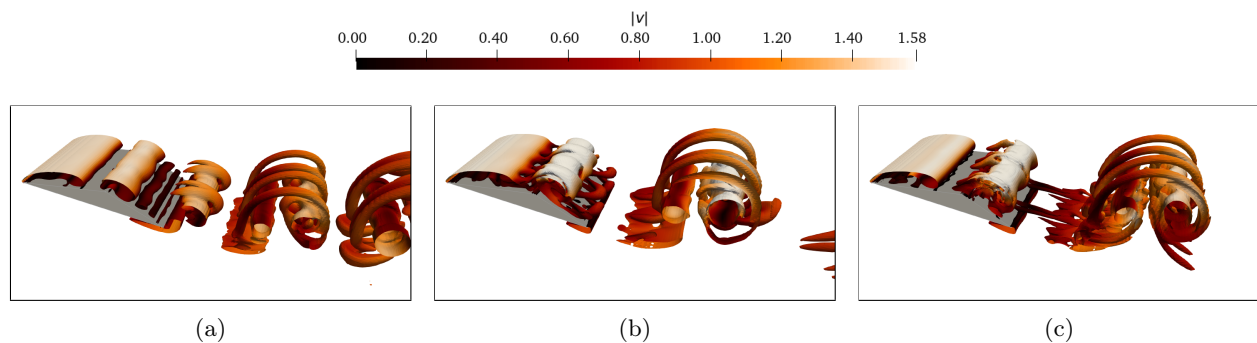


Figure 14: Images of instantaneous Q-criterion isosurfaces  $Q = 1$ , coloured by velocity magnitude  $|\mathbf{v}|$ , for triangular airfoils with minimum  $C_D$  (a), high  $C_L/C_D$  (b), and maximum  $C_L$  and  $C_L/C_D$  (c) obtained from 3D DNS optimization at generation 15 for  $\alpha = 12^\circ$ .

Figure 13 shows the time-averaged velocity magnitude in the midspan plane with superimposed velocity line integral convolutions (LICs) for three optimum airfoil shapes found by the  $\alpha = 12$  degree 3D DNS shape optimization. Figure 14 shows images of instantaneous Q-criterion isosurfaces  $Q = 1$ , coloured by velocity magnitude over the same three optimum airfoil shapes. The three optimum shapes correspond to the 3D DNS optimum airfoils represented with lines in Figure 12b. The time-averaged LICs in Figure 13 show that, for the three cases, the flow separates from the leading edge, with vortices attached to the upper surface of the airfoil. Figure 14 shows the unsteady lift-generating physics for the three examined airfoil shapes, with three-dimensional breakdown of the stream-wise vortices in the span-wise direction. As investigated previously by Caros et al. [7], the results obtained with 2D DNS show higher lift because they do not resolve the three-dimensional breakdown of the stream-wise vortices that leads to less coherent vortices rolling on

the suction side of the airfoil.

As the examined airfoils inhabit the extremes in the objective space, it may be presumed that all the airfoils found in the Pareto front experience similar separation phenomena with three-dimensional flow structures breaking the coherence of the stream-wise vortices. When optimizing at  $\alpha = 12^\circ$ , 3D DNS should be used to capture the fundamentally three-dimensional nature of the lift-generating flow physics.

## 5 Conclusion and Future Work

Triangular airfoil shapes with sharp leading edges have been optimized for Martian atmospheric conditions, targeting aerodynamic performance. Optimizations have been performed with both 2D and 3D DNS with PyFR, at angles of attack  $\alpha = 6^\circ$  and  $\alpha = 12^\circ$ .

For low angles of attack ( $\alpha = 6^\circ$ ), it has been shown that the optimal airfoils induce predominantly two-dimensional flow. 2D and 3D DNS evaluations have given similar aerodynamic coefficients, which resulted in similar optimum airfoil shapes for both optimizations. Furthermore, the 2D DNS optimum airfoils evaluated with 3D DNS are still within the 3D DNS Pareto optimal set of solutions. Thus, 2D DNS optimization can be used to optimize triangular airfoils at low angles of attack such as  $\alpha = 6^\circ$ , under the conditions tested.

On the other hand, at an angle of attack of  $\alpha = 12^\circ$ , optimization with 2D and 3D DNS has given very different results throughout the optimization. It has been shown that 2D DNS evaluations over-predict lift and under-predict drag at  $\alpha = 12^\circ$  due to the lack of three-dimensional vortex-breakdown physics. When evaluated with 3D DNS, the 2D DNS optimum airfoils fall into the dominated area of the 3D DNS Pareto front. Hence, 3D DNS optimization is necessary to optimize triangular airfoils at high angles of attack such as  $\alpha = 12^\circ$ , under the conditions tested.

Future work should investigate the use of 2D DNS optimization results as initial sampling for 3D DNS optimization at high angles of attack. This strategy might accelerate the optimization convergence since the optimum airfoils from the 2D DNS optimization are closer to the optimal ones than the original Latin Hypercube Sampling. Also, convergence acceleration strategies should be investigated to reduce the cost of 3D DNS optimization, such as the use of surrogate models.

## 6 Acknowledgements

The authors are grateful for support from the Engineering and Physical Sciences Research Council via an EPSRC Fellowship (EP/R030340/1) and for the granted allocation on the Piz Daint supercomputer at the Swiss National Supercomputing Center (CSCS) which has made this work possible.

## References

- [1] J. Bob Balaram, Timothy Canham, Courtney Duncan, Matt Golombek, Håvard Fjær Grip, Wayne Johnson, Justin Maki, Amelia Quon, Ryan Stern, and David Zhu. Mars helicopter technology demonstrator. *AIAA Atmospheric Flight Mechanics Conference*, 2018.
- [2] Witold J. F. Koning, W. Johnson, and H. F. Grip. Improved Mars helicopter aerodynamic rotor model for comprehensive analyses. *AIAA Journal*, 57:3969–3979, 2019.
- [3] Phillip Munday, Kunihiko Taira, Tetsuya Suwa, Daiju Numata, and Keisuke Asai. Non-linear lift on a triangular airfoil in low-Reynolds-number compressible flow. *Journal of Aircraft*, 52:924–931, 2015.
- [4] Witold J.F. Koning, Ethan A. Romander, and Wayne Johnson. Performance optimization of plate airfoils for Martian rotor applications using a genetic algorithm. *45th European Rotorcraft Forum 2019, ERF 2019*, 1, 2019.
- [5] Witold J. F. Koning, Ethan A. Romander, and Wayne Johnson. Optimization of low Reynolds number airfoils for Martian rotor applications using an evolutionary algorithm. *AIAA Science and Technology Forum and Exposition (AIAA SciTech)*, 2020.
- [6] Gaku Sasaki, Tomoaki Tatsukawa, Taku Nonomura, Akira Oyama, Takaaki Matsumoto, and Koichi Yonemoto. Multi-objective optimization of airfoil for Mars exploration aircraft using genetic algorithm.



*Transactions of the Japan Society for Aeronautical and Space Sciences, Aerospace Technology Japan*, 12:59–64, 2014.

- [7] L. Caros, O.R.H. Buxton, T. Shigeta, T. Nagata, T. Nonomura, K. Asai, and P.E. Vincent. Direct numerical simulation of flow over a triangular airfoil under Martian conditions. *AIAA Journal*, March 2022.
- [8] F. D. Witherden, A. M. Farrington, and P. E. Vincent. PyFR: An open source framework for solving advection-diffusion type problems on streaming architectures using the flux reconstruction approach. *Computer Physics Communications*, 185(11):3028–3040, 2014.
- [9] H.T. Huynh. A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods. *AIAA Paper 4079*, pages 1–42, 2007.
- [10] Christopher A. Kennedy, Mark H. Carpenter, and R. Michael Lewis. Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations. *Applied Numerical Mathematics*, 35(3):177–219, 2000.
- [11] D. Goldberg. *Genetic Algorithm in Search, Optimization, and Machine Learning*, volume XIII. Addison-Wesley Longman Publishing Co., Inc., 1989.
- [12] D. Kalyanmoy, A. Sameer, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization NSGA-II. *IEEE Transaction on Evolutionary Computation*, 6:182–197, 2002.
- [13] Julian Blank and Kalyan Deb. Pymoo: multi-objective optimization in Python. *IEEE Access*, PP:1–1, 04 2020.
- [14] Eckart Zitzler and Lothar Thiele. Multiobjective optimization using evolutionary algorithms - a comparative case study. *Lecture Notes in Computer Science*, 1498 LNCS, 1998.
- [15] L. Caros, J. Blank, O.R.H. Buxton, and P.E. Vincent. Comparing strategies for DNS based optimization of airfoils for Martian rotorcraft. *Vertical Flight Society Forum* 78, May 2022.