JOURNAL OF AIRCRAFT
Vol. 48, No. 1, January–February 2011

Robust Design Optimization of Airfoils with Respect to Ice Accretion

Tiziano Ghisu,* Jerome P. Jarrett,† and Geoffrey T. Parks‡
University of Cambridge, Cambridge, England CB2 1PZ, United Kingdom

DOI: 10.2514/1.C031100

Ice formation on an aircraft happens when supercooled water droplets impinge on the leading-edge surfaces, freezing immediately upon impact. Ice accretions alter the effective shape of the aircraft, modifying the aerodynamic forces and moments caused by the air flow over iced components. The primary effects are increased drag, reduced stall angle, and reduced maximum lift. Additional effects include decreased pitching and hinge moments that lead to a reduction in the effectiveness of control surfaces. Over the last 80 years, significant efforts have been devoted to improving understanding both of the physics of the ice accretion process and of the resulting degradation in aerodynamic performance. Significant progress has also been made through improved certification processes and more effective ice-protection systems. Despite these continued efforts, icing-induced incidents and accidents continue to occur on all classes of aircraft, to the point that the National Transportation Safety Board ranked icing among its "Most Wanted Aviation Transportation Safety Improvements." To take full advantage of these improved predictive capabilities, they need to become an active part of the design process. In this study, a computational fluid dynamics analysis of the degradation in airfoil performance caused by a quarter-round ice shape (simulating the presence of a supercooled large droplet ice accretion or runback ice) has been included in an optimization system featuring a metaheuristic multi-objective optimizer and a geometry modeler based on a freeform deformation of a datum geometry. In recognition of the real-world variability in the ice accretion location, this parameter has been treated as uncertain, with a given probability density function. An adaptive nonintrusive polynomial chaos approach has been used to propagate this input uncertainty through the otherwise deterministic system. The results presented show how, using the computational design approach developed, considerable improvements in the iced performance of the airfoil can be achieved over a range of ice locations while still maintaining good clean-wing characteristics.

Nomenclature

C_d	=	drag coefficient
C_l	=	lift coefficient
C_m	=	pitching moment coefficient
C_p	=	pressure coefficient
c	=	airfoil chord
I_i	=	generic Wiener–Askey polynomial
L	=	ice accretion location
l, m, n	=	dimensions of the freeform deformation lattice
N_0	=	quadrature order
P_i^*	=	adaptive orthonormal polynomial
$\mathbf{P}_{ijk}$	=	new position of a freeform deformation control point
p	=	polynomial chaos order
r_{LE}	=	leading-edge radius
s, t, u	=	coordinates of a generic point before freeform deformation
t/c	=	thickness-to-chord
W	=	weighting function
w_i	=	weight factor
$\mathbf{X}_{\mathrm{FFD}}$	=	coordinates of a generic point after freeform deformation
x	=	position along chord
Y	=	random process
$y_i, c_i, \hat{c}_i$	=	expansion coefficients
α	=	angle of attack
β_{TE}	=	trailing-edge angle
ΔC_l	=	change in lift coefficient due to ice accretion
θ	=	generic random event
μ	=	mean
ξ_i	=	support random variable
σ	=	standard deviation
Φ_i	=	Wiener–Askey chaos

Received 7 May 2010; accepted for publication 16 July 2010.

*Research Associate, Engineering Design Center, Engineering Department, Trumpington Street.

†University Lecturer, Engineering Design Center, Engineering Department, Trumpington Street. Member AIAA.

‡Senior Lecturer, Engineering Design Center, Engineering Department, Trumpington Street.

I. Introduction

IN 1940, Clarence L. Johnson of Lockheed Aircraft Corporation (now Lockheed Martin) described the problem of ice accretion on aircraft surfaces as "one of the most important ones facing the aviation industry" [1]. Almost 70 years later, a Bombardier DHC8-402 Q400 operating as Flight 3407 went through severe pitch and roll oscillations, lost altitude and then crashed into a private home five miles short of the runway in the northeast Buffalo suburb of Clarence Center. Despite the anti-icing system being turned on 11 minutes into the flight, the crew were discussing the consistent presence of ice on the aircraft's wings and windscreen just minutes before the crash. All 45 passengers, the four members of the air crew and one person in the house were killed [2]. In the last 80 years, despite significant resources being devoted to improving understanding of the ice accretion process and to the design of better methods of prevention, incidents and accidents caused by icing-induced loss of control continue to occur on all classes of aircraft. The problem of icing has been recognized by many organizations: for instance, the National Transportation Safety Board (NTSB) has included it among its "Most Wanted Aviation Transportation Safety Improvements" since 1997 [3].

Ice accretion on an aircraft occurs when supercooled water droplets (droplets of liquid water below its freezing temperature) impinge on the aircraft, freezing immediately upon impact. The type of accretion depends on the air temperature, liquid water content and size of the droplets: low temperatures, low water content and small

droplets lead to immediate freezing, producing a brittle, opaque accretion (rime ice). Larger and warmer droplets tend to deform and flow on the surface before freezing, generating a clear accretion with a smoother surface (glaze ice). Glaze ice is more difficult to detect and more dangerous as it tends to create protuberances (horn-type accretions) which can be particularly detrimental to aircraft performance. Aircraft certification requires them to fly in a range of natural icing conditions with no significant effects on their performance. The Federal Aviation Regulations (FAR) (Part 25 Appendix C) define icing envelopes (in terms of air temperature, liquid water content and droplets' size) for aircraft certification, corresponding to 99.9% of icing conditions found in stratiform clouds [4]. Most light aircraft are not required to pass this certification and do not usually have the required ice-protection systems for flying into icing conditions. Larger aircraft are equipped with a variety of anti-icing or de-icing systems (heaters, pneumatic boots and liquid flows) to help them prevent ice formation [5].

The crash of an ATR-72 (American Eagle Flight 4184) on 31 October 1994 led to a complete reevaluation of the understanding of aircraft icing. The accident occurred due to freezing drizzle, consisting of droplet sizes much larger than those prescribed by the FAR [6]. Supercooled large droplets (SLD) can form in several ways: the penetration of a warm front through a cold layer of air causing snow to melt and then cool rapidly without freezing, and the presence of droplets falling at different speeds which collide and coalesce into larger drops are common formation patterns. Larger droplets have a greater inertia which moves their impingement location further back on the airfoil or wing surface. SLD icing conditions can result in the formation of ice downstream of the de-icing boot, leading to a large separation bubble which grows with increasing angle of attack until it fails to reattach, causing premature stall. The result is large losses in lift, increases in drag and changes in pitching moment and aileron hinge moment, which can cause a change in the aileron stick force leading to hinge moment reversal and aileron snatch [7]. However, the problem is more general, because similar accretions can also occur in the presence of supercooled droplets of the size described by the FAR. In this case, the accretion occurs behind the boot because of surface water runback (water not freezing completely on a pneumatic-boot or not being evaporated fully by a heated leading edge) [6].

The ATR-72 crash reinvigorated research efforts in this area, moving the focus to the effects of spanwise ridges (typical of SLD conditions and runback ice accretions) on aerodynamic performance and to the sensitivity of different airfoil shapes. The two fundamental questions asked in this research are: 1) where and how will the ice shape form? and 2) what will the consequences on aerodynamic performance be? These two questions have been usually treated separately due to the complexity of the phenomena involved [8]. Experimental (icing tunnels for accretion studies and traditional wind tunnels with simulated ice shapes to study the effects on the airfoil/wing performance) and computational [analytical ice accretion codes, such as LEWICE and GlennICE [9], or computational fluid dynamics (CFD) for performance analysis] approaches have both been used, improving understanding of the phenomena involved and of the changes required to reduce the sensitivity of aircraft to ice accretion.

While the capabilities to correctly reproduce (both experimentally and numerically) the formation of ice protuberances on a wing or other surface and to predict their impact on aerodynamic performance are fundamental, they do not in themselves produce a better design. To take full advantage of these improved predictive capabilities, they need to become an active part of the design process [10]. Design optimization is an expanding field: improved solutions to a design problem can be found effectively by linking an optimizer to appropriate evaluation tools, while also reducing design times through automation and providing the designers more time for creative thinking [11]. The key enabler for effective exploration of the design alternatives is the parameterization scheme, which allows the design to be modified in the quest for a better solution.

A common simplification in the design of engineering systems is to assume nominal values for a number of design parameters. While such assumptions are often essential to simplify the design problem and reduce design times to the point where solution of the problem is a practicable proposition, they may lead to products that exhibit significant deviations from their nominal performance in real-world conditions. This problem is often exacerbated in heavily optimized products, which tend to lie in extreme regions of the design space [12]. The introduction of some consideration of robustness is essential to guarantee that the results of the optimization process can maintain the expected performance improvements in real-world conditions.

In this work, an open-source CFD solver has been used to analyze the degradation in performance due to the presence of a quarter-round ice accretion (typical of SLD or runback ice conditions) on an airfoil. This analysis method, validated by comparison with existing experimental work, was integrated within an optimization system featuring a metaheuristic optimizer and a geometry modeler based on a freeform deformation (FFD) of an initial geometry. In recognition of the fact that clean and iced airfoil performance represent not necessarily commensurable design objectives, a multi-objective approach was adopted: seeking simultaneous improvements in both clean and iced performance, the optimizer aims to find the tradeoff between these potentially conflicting figures of merit, in the form of a number of equivalent (nondominated) design solutions, for the decision-maker to choose from.

Further, in recognition of the importance yet uncertainty of the ice accretion location on airfoil performance, this parameter was treated as uncertain but with a given probability density function (PDF). An adaptive nonintrusive polynomial chaos (NIPC) approach (based on a set of orthonormal polynomials defined in real time with respect to the input PDF [13]) facilitates an efficient analysis of the performance variability in the presence of the given uncertain parameter while using the same deterministic solver. The robust optimization process leads to the identification of a number of design solutions with improved iced performance over a range of possible ice locations, with performance that is still acceptable in clean conditions.

The rest of the paper is structured as follows. Section II reviews past work on ice formation and its effect on aircraft performance. Section III describes the computational approach to the evaluation of clean and iced airfoil performance used in this study and presents validation information to confirm its fitness for purpose. Section IV presents details of the design optimization system used in this work, while Sec. V describes the polynomial chaos (PC) approach used to handle uncertainty to facilitate robust optimization. Section VI presents and discusses the results from a number of studies seeking to optimize iced airfoil performance using the design optimization system developed, and finally Sec. VII draws conclusions from this work.

II. Background

One of the first studies of the problem of ice formation on aircraft is due to Carroll and McAvoy [14] in 1929. Their work concentrates on analyzing the atmospheric conditions under which ice is deposited on the exposed surfaces of aircraft in flight. They identify the two main types of ice formation (rime and glaze ice) and the main effects on aircraft performance (loss of lift and increase in drag being more serious than the increase in weight due to the ice). Their main recommendation is the avoidance of conditions under which ice formation is most likely to occur.

Jacobs [15] studies the effects of protuberances on the performance of a NACA 0012 airfoil, with particular attention to the consequences of their position, size and shape. He concludes that protuberances larger than 0.5% of the airfoil chord can cause a severe drag penalty and that, while the effect on lift of a protuberance on the lower surface is unimportant, it can become significant for a protuberance on the upper surface as its position approaches the leading edge, even for a protuberance so small that it would ordinarily be classified as a surface roughness. Noth and Polte [16] identify the greater danger of SLD conditions and suggest their avoidance, relying on better meteorological predictions rather than less sensitive

aircraft or effective anti-icing systems. Gulick [17] recognizes the importance of de-icing systems and studies the effects of the types of ice accretions that remain on a wing in spite of the operation of pneumatic de-icers (irregular ridges on both sides of the inflatable rubber boots), on a NACA 0012 airfoil. He observes a reduction in the lift curve slope at all useful angles of attack, a reduction in the maximum lift coefficient from 1.32 to 0.80 and a 90% increase in drag. He also identifies the loss in lift (and the consequent increase in stalling speed) as the main concern for the aircraft's safe operation. The effects of spanwise ridge accretions are also studied by Johnson [1], on a Lockheed Electra airplane, measuring a 32% reduction in lift, a 47% increase in drag and a significant reduction in angle of attack for maximum lift, no change in longitudinal stability but a 64% reduction in the available rolling moment. Gray and Von Glahn [18] study the effects of primary, runback ice and frost formation on a NACA 65-212 airfoil for different icing conditions, identifying the higher risks of glaze ice accretions (both primary and runback) over rime ice formations occurring at lower temperatures.

The importance of a better understanding of the change in performance characteristics due to ice accretion, together with the need to test new de-icing and anti-icing systems, convinced NACA to build an icing research tunnel (IRT) at the Lewis Flight Propulsion Laboratory in Cleveland, Ohio [19]. Work at IRT led to the publication of empirical correlations for the change in drag coefficient caused by ice formations on a NACA 65A004 airfoil [20]. The effects of icing time, airspeed, air total temperature, liquid water content, cloud droplet impingement efficiencies, airfoil chord length and angles of attack were considered in this work. Gray [21] later extended the correlations including available data for other airfoils with thickness ratios up to 15%, and also presented a correlation for the changes in lift and pitching moment coefficients due to ice on a NACA 0012 airfoil.

After a period of relatively little icing research, interest in the subject has renewed at end of the 1970s, due to a combination of factors: the development of new classes of aircraft, the characteristics of which had not yet been studied under icing conditions; the development of new ice-protection systems; and the recent advances in CFD, which also pushed experimental research to produce the data required for validation [19]. One of the first numerical studies is due to Potapczuk and Gerhart [22], who solve the Reynolds-averaged Navier–Stokes (RANS) equations in a body-fitted curvilinear coordinate system for a NACA 0012 airfoil in the presence of glaze- and rime ice accretions. Kwan and Sankar [23] use a three-dimensional, compressible RANS solver to predict the performance of an untapered and untwisted wing with a NACA 0012 planform shape and an aspect ratio of five, numerically demonstrating the decrease in the wing stall angle as a result of leading-edge ice formation. Potapczuk and Berkovitz [24] analyze experimentally both ice accretion patterns and performance degradation of a multi-element airfoil under different configurations, in the presence of glaze, rime and mixed ice conditions. They also compare the experimental shapes with the predictions obtained from LEWICE, an analytical ice accretion prediction code developed at NASA Lewis Research Center, and consisting of a two-dimensional potential flow panel method to calculate the flowfield, a particle trajectory code to calculate the water droplet impingement patterns, and an ice accretion module that solves the quasi-steady energy balance equation, obtaining better results for rime ice than for glaze ice accretions. They also use a RANS solver to predict iced airfoil performance, obtaining poor agreement due to the coarseness of the grid. Both numerical predictions were performed for a single-element airfoil.

Most of the icing studies before 1994 concentrate on symmetrical airfoils (often a NACA 0012) with large glaze ice shapes. The American Eagle Flight 4184 accident changed the focus of icing research: loss of this ATR-72, thought to have been caused by an SLD ice accretion, drew attention to different ice shapes and to the sensitivity of different airfoil sections [19]. Renewed efforts were devoted to the determination of critical ice shapes and their effects on aerodynamic performance. Addy et al. [25] study the effects of ice accretions in SLD conditions on a NACA 23012 airfoil operating a de-icing boot, including the effects of temperature, droplet size, angle of attack, airspeed and boot-cycle time. They report the formation of a ridge aft of the active portion of the de-icing boot under nearly all operating conditions, with the location and shape depending on the test conditions. Lee and Bragg [26] conduct an experimental study to understand the effects of simulated spanwise-step ice shapes (forward- and backward-facing quarter-rounds, half-round and forward-facing ramp) on airfoil aerodynamics, on NACA 23012 and NLF 0414 airfoils. They observe large differences in the aerodynamic penalties (maximum C_l as low as 0.25 for the former, 0.68 for the latter), a strong dependence on the ice location and a mild dependence on the ice shape. They attribute the different behavior to the different pressure distributions of the two airfoils: the NACA 23012 is a front loaded airfoil with a pronounced leading-edge suction peak, while the NLF 0414 has a relatively uniform chordwise loading until about 70% chord.

Dunn et al. [6] conduct a numerical study on the effects of simulated spanwise-step ice accretions, resulting from SLD icing conditions, on a NACA 23012 airfoil with a simple flap. Using an adaptive unstructured mesh, they correctly predict the trends for all the main aerodynamic features associated with the presence of ice of the defined shape. They examine the sensitivity of the results to the ice shape size and conclude that significant reductions in lift can be obtained even for relatively small protuberances (0.83 to 1.39% of the chord), a conclusion consistent with previous experimental findings. Chung et al. [27] perform a computational study on the effects of a ridge-ice formation on a NACA 23012 airfoil and on a three-dimensional wing. Their work aims at providing some qualitative trends and insights into the aerodynamics that led to a control system problem in a commercial turboprop-powered aircraft accident under investigation by the NTSB. The ice shape was selected because it was the one that caused the largest lift decrease and drag increase among those analyzed in a preliminary study. The trends for the two-dimensional and three-dimensional simulations are similar, although the maximum lift for the latter occurs 4 degrees later than for the two-dimensional case. Kumar and Loth [28] investigate the effects of an upper surface spanwise ice shape on the aerodynamic performance of a NACA 23012 and several airfoils from the NASA Modern Airfoil program. They compare the numerical results with available experimental data, for different values for ice location, size and shape, obtaining good predictions for the main aerodynamic forces and moments. The presence of the ice shape produces a large separation bubble, grossly changing the pressure distribution and modifying the airfoil from leading-edge to thin-airfoil stall. The lift prediction deteriorates once the upper surface flow is fully separated, because of the inability to capture the highly unsteady vortex-shedding at this condition. The authors note how the NACA 23012 exhibits the largest performance degradation, with a critical ice location close to 10% chord. Pan and Loth [29] extend the previous work considering the effects of Reynolds and Mach numbers, which they find to be considerably smaller than on clean airfoils. They investigate the critical ice location for different airfoils, noting how this tends to be between the positions of maximum pressure gradient and minimum pressure, but also how an airfoil with a larger peak suction does not always suffer more severe performance degradation under icing conditions. In addition, they analyze the performance degradation of a NACA 23012 tapered wing, which presents a lower lift curve slope (due to its finite span) and a delayed stall, but also a stronger break in the lift curve than the corresponding airfoil. The same authors [30] use a detached eddy simulation (DES) approach to improve the predictive capabilities of the stall behavior for iced airfoils, obtaining more satisfactory results than with the RANS approach: the DES results are able to qualitatively capture the flat pressure plateau after the ice shape and the stall behavior (where the RANS simulations only predicted a break in the lift curve slope). Marongiu et al. [8] compare the numerical predictions of three different flow solvers, with two turbulence models each, for a NLF 0414 airfoil with leading-edge ice shape. They observe significant discrepancies in the predictions from the different codes in the poststall region, attributable to their different dissipation properties, but also find that the Spalart–Allmaras

model produces results closer to the experimental data than the k-ω shear stress transport model.

III. Computational Approach

OpenFOAM, an object-oriented open-source package for numerical simulation in continuum mechanics, written in C++, was chosen for the numerical simulation of airfoil performance. OpenFOAM consists of a series of field operation and manipulation libraries, which are used to create executables (or applications). The applications fall into two categories: solvers (each designed to solve a specific problem in continuum mechanics) and utilities (designed to perform tasks that involve data manipulation). The OpenFOAM distribution contains numerous solvers and utilities covering a wide range of problems, while additional ones can be created by users with knowledge of the underlying method, physics and programming techniques involved [31].

The blockMesh utility was used to create geometry and mesh: it allows the generation of block-structured meshes by dividing the domain into a number of hexahedral blocks, the edges of which can be straight lines, arcs or splines. The mesh is specified through the number of cells and growth rate in each direction, for each block. Figure 1a shows a view of the block structure for the fluid domain, while Fig. 1b presents a closer view of the blocks near the airfoil. The inner block outer limits were obtained by means of a FFD of the airfoil, in order to have a grid as orthogonal as possible near the airfoil, in the boundary layer region, where the mesh quality is particularly important. The domain extends 9 chords upstream and 10 chords downstream of the airfoil.

A grid-independence study was conducted by calculating the variation of the main aerodynamic coefficients (lift and drag) with the number of grid points in the different blocks (leading-edge, central, trailing-edge, wake regions), in the directions both parallel and perpendicular to the airfoil chord. The mesh chosen for the clean airfoil simulations consists of 99,000 cells and is shown in Fig. 2a, while Fig. 2b presents a closer view of the mesh near the airfoil. A $y+$ value close to unity was used to allow a good prediction of the boundary layer in the presence of separated regions.

The solver rhoSimpleFoam was used to simulate the flow around the airfoil: it solves the RANS equations using a pressure correction (SIMPLE) algorithm, allowing different numerical schemes and turbulence models to be readily substituted, in keeping with the OpenFOAM philosophy. The GammaV scheme was chosen for approximating the convective terms: this is an improved version of the Gamma scheme formulated to take into account the direction of the flow field. The Gamma scheme is a smooth and bounded blend between the second-order central differencing scheme and the first-order upwind scheme (the former is used wherever it satisfies the boundedness requirements, the latter otherwise) [32]. Central differencing was used for the diffusive terms and a Gauss–Siedel approach to solve the linearized systems of equations, with a geometric agglomerated algebraic multigrid preconditioner for convergence acceleration. The Spalart–Allmaras model was used for turbulence closure, as it has led to satisfactory results in a number of previous studies [8,29].

Figure 3 presents a comparison between the C_l-α and C_d-α curves obtained with rhoSimpleFoam and the experimental results of Broeren et al [33], for a NACA 23012 airfoil at a free stream Mach number of 0.18 and Reynolds number of 1.8×10^6. The agreement is satisfactory: the maximum C_l is well predicted, while the C_d at high angles of attack is slightly overpredicted.

The effects of a spanwise ridge (simulating the presence of SLD ice accretions) was analyzed using a quarter-round ice shape. The numerical grid for the iced conditions was obtained by superimposing the ice shape on the grid for the clean airfoil case and running the snappyHexMesh utility, which generates meshes from triangulated surface geometries in stereolithography format, conforming to the surface by iteratively refining a starting mesh and morphing the resulting mesh to the surface [31]. The resulting grid, containing 103,660 cells, is shown in Fig. 4.

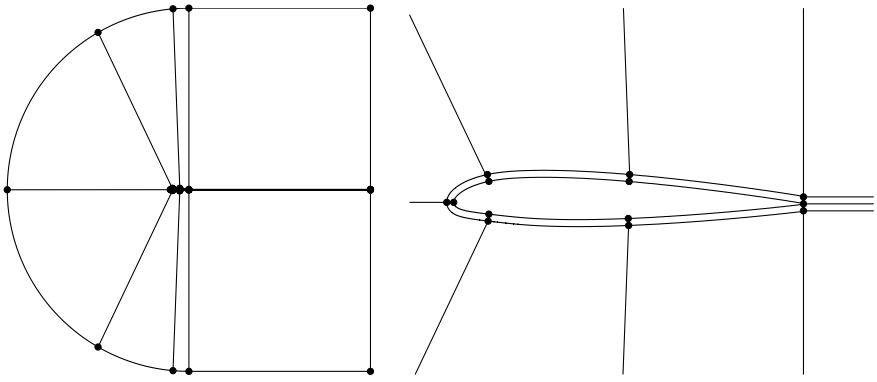

a) Overall domain **b) Close-up view of the airfoil**

Fig. 1 Structure for the multiblock mesh.

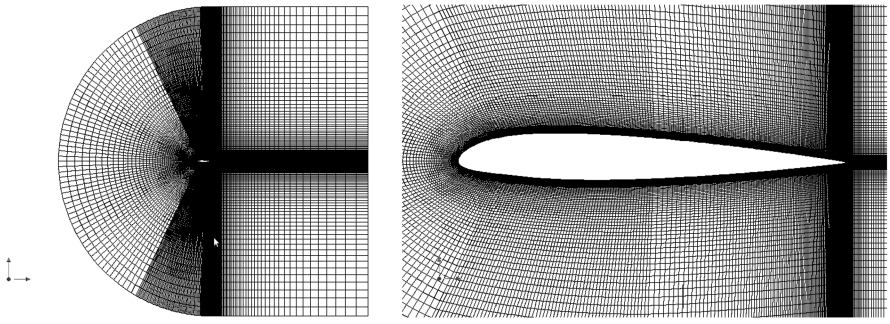

a) Overall domain **b) Close-up view of the airfoil**

Fig. 2 Computational grid around the airfoil.

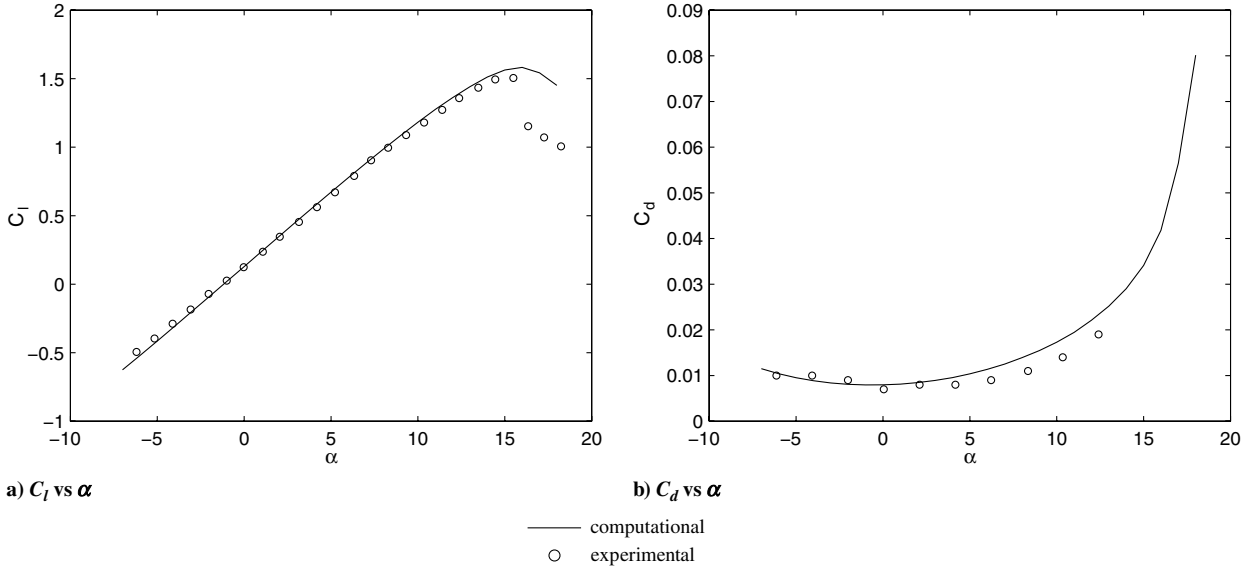

Fig. 3 Lift and drag coefficient variation with angle of attack.

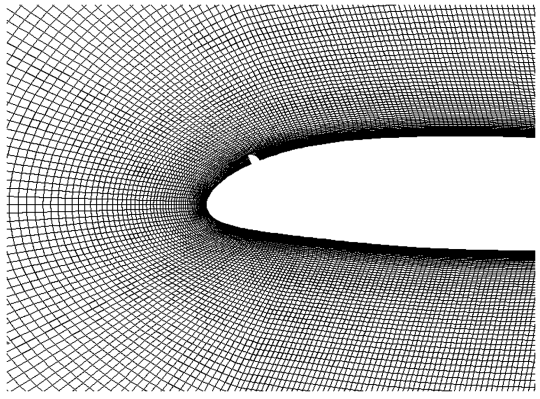

Fig. 4 Computational grid for the iced airfoil.

Figure 5 presents a comparison between the experimental results of Broeren et al. [33] and the predictions obtained with OpenFOAM, for a 1.39% chord quarter-round ice shape located at 10% chord, for the same Mach and Reynolds numbers as in the clean simulations. The same comparison is also presented for the same ice shape located at 20% chord (Fig. 6).

At low angles of attack, OpenFOAM gives good agreement with the experimental data, although the slope of the lift curve is slightly underpredicted. The solver fails to predict the leveling off of the lift curve (at about 3°), while it predicts a break in the lift curve slope a few degrees earlier. It correctly predicts the further increase in lift at higher angles of attack, due to lift still being generated from the portion of the airfoil upstream of the ice shape [33]. The computational results match closely the experimental trends for low angles of attack (when a small separation bubble is generated and reattaches soon after the ice accretion) and for large angles of attack, in the presence of a fully separated flow. The largest differences are in between these cases, in the region of rapid nonlinear growth of the recirculation bubble. Comparable results have been obtained by [6,8,28,34]. Pan and Loth [29] attribute these differences to the inability of the RANS approach to capture the highly unsteady separation region developing before stall. In [30], the same authors satisfactorily predict the plateau in the lift curve with a DES approach (albeit with an underprediction in the value for the maximum C_l). The predictions of C_d match more closely the experimental results.

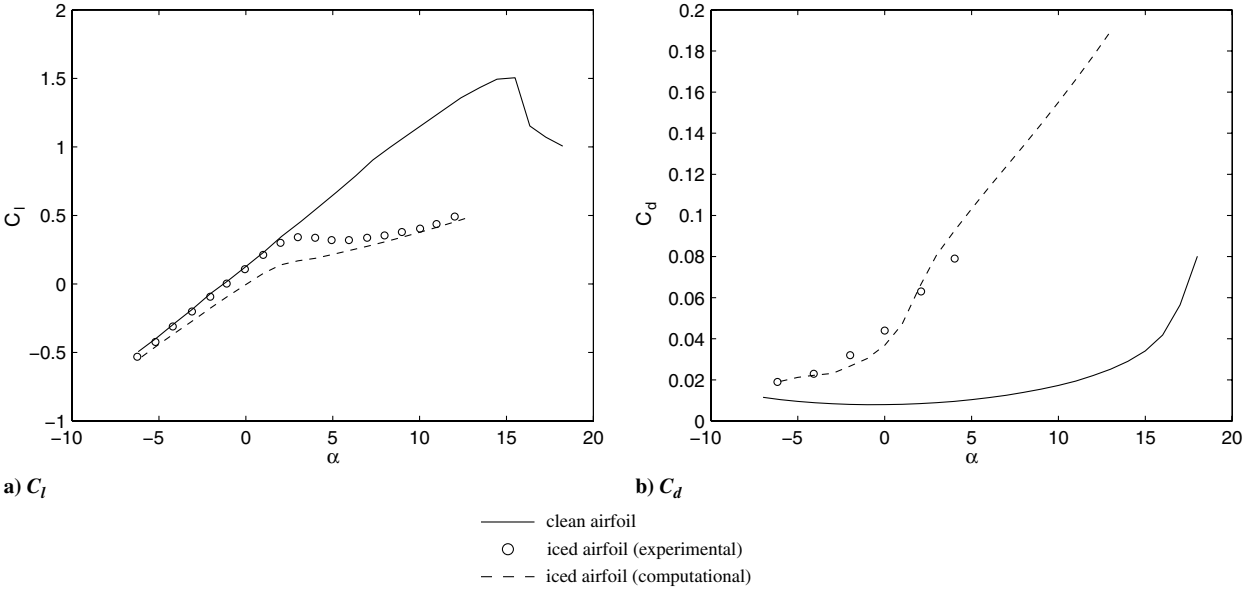

Fig. 5 Effect of 1.39% chord quarter-round ice at 10% chord.

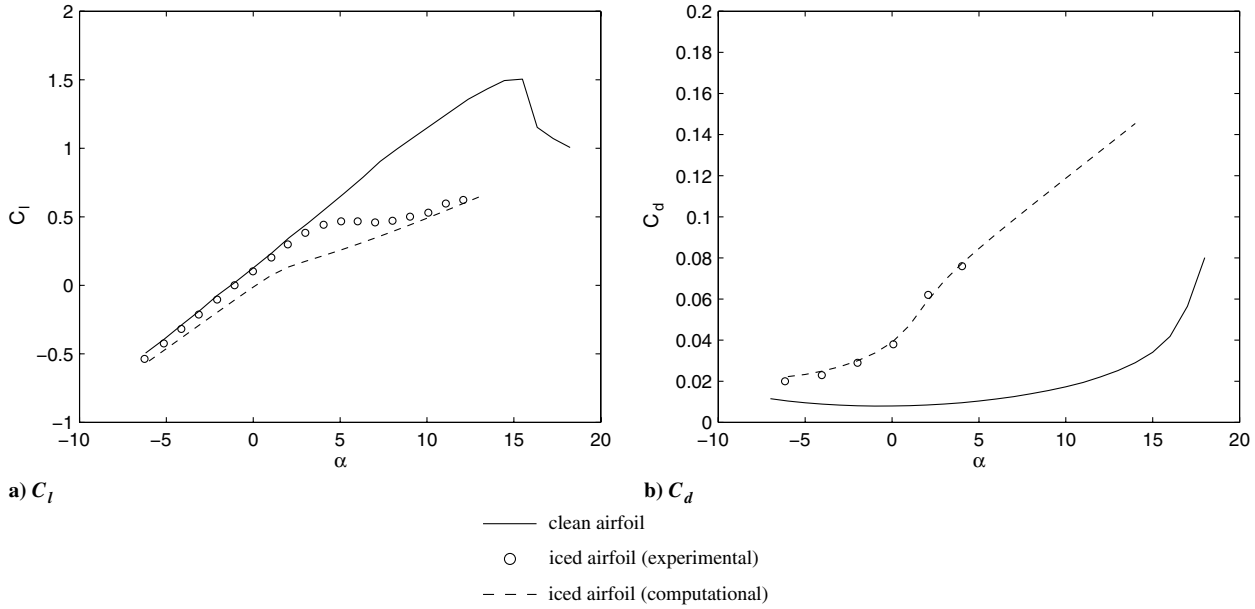

Fig. 6 Effect of 1.39% chord quarter-round ice at 20% chord.

To verify the independence of the numerical results on the grid used for these simulations, this was modified by doubling the number of grid nodes in the directions either parallel (refinement 1) or perpendicular (refinement 2) to the airfoil chord. The results (presented in Fig. 7) justify the choice of the grid.

Figure 8 compares the pressure coefficient along the airfoil surface at different angles of attack with experimental data [34], for a 0.83% chord quarter-round ice shape located at 10% chord. OpenFOAM accurately predicts the pressure variation over the majority of the airfoil surface, giving excellent agreement in nonseparated regions (pressure surface and suction surface before the ice shape). The pressure increase in the stagnation region in front of the ice shape is also well predicted as is the magnitude of the suction peak downstream of the quarter-round ice shape. However, the computed pressure recovery is faster than that observed experimentally, with the discrepancy increasing with the size of the separation bubble. This can be ascribed to the inadequacy of the turbulence model in predicting the amount of entrainment within the boundary layer.

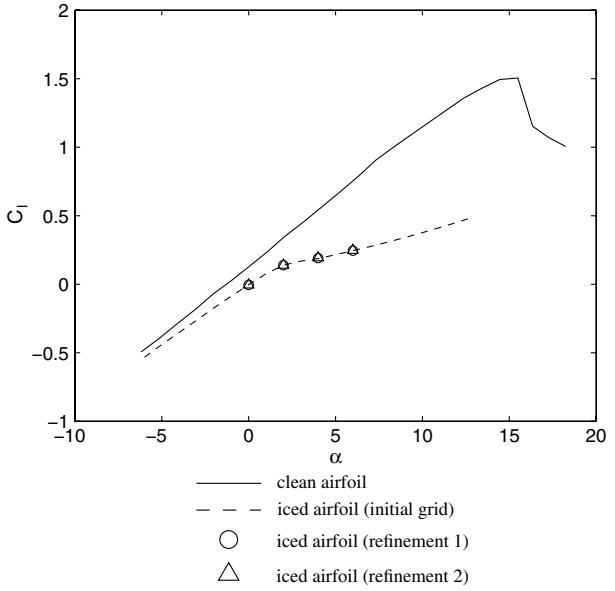

Fig. 7 Verification of the chosen computational grid for a 1.39% chord quarter-round ice at 10% chord.

Figure 9 compares the pressure coefficient along the airfoil surface for different sizes of the quarter-round ice shape with experimental data [34], for an angle of attack of 3° and with the ice accretion located at 10% chord. Despite the difficulties in capturing the size of the separation bubble, the trends are correctly predicted.

Figure 10 shows the dependence of the pressure coefficient along the airfoil surface on the ice shape position, for a 0.83% chord

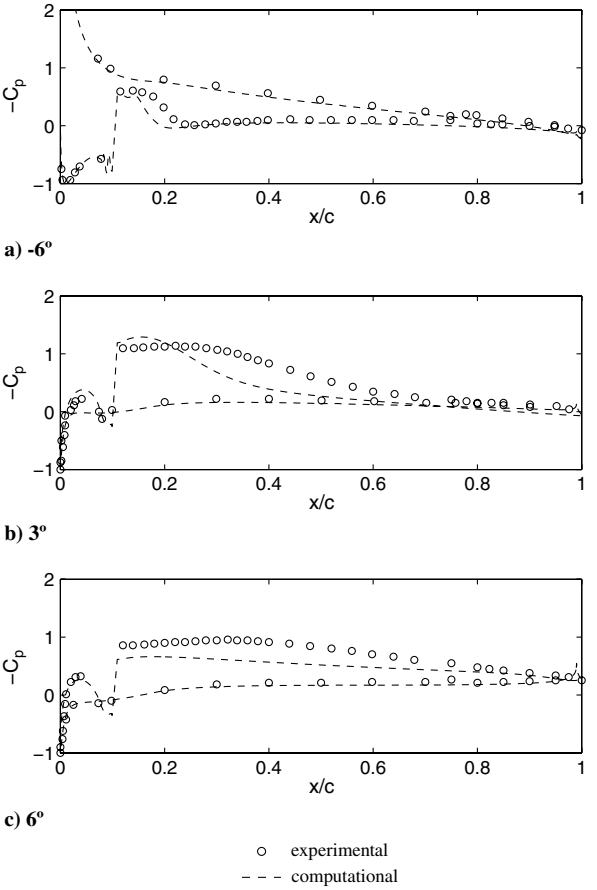

Fig. 8 Comparison of C_p distributions at different angles of attack, for a NACA 23012 airfoil with a 0.83% chord quarter-round ice accretion at 10% chord.

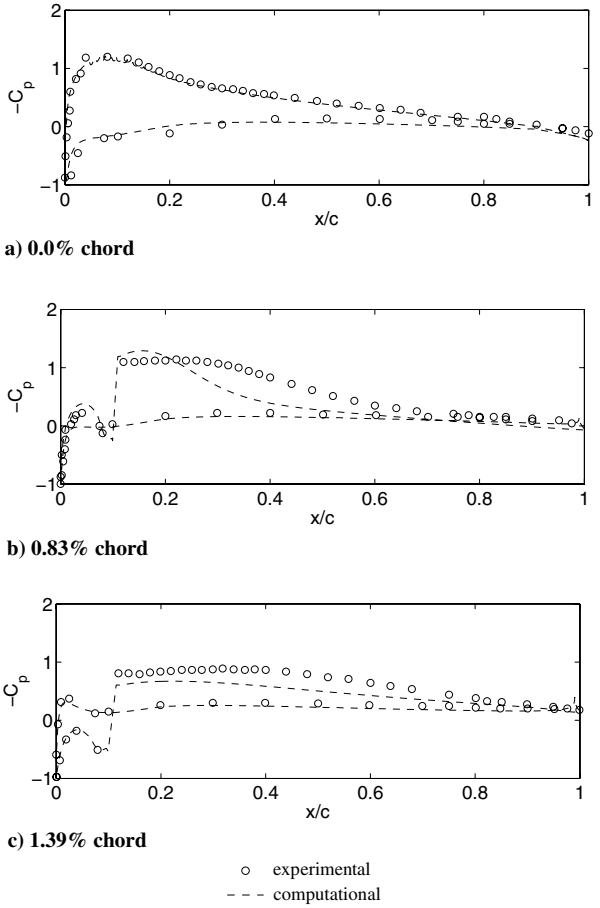

Fig. 9 Effect of ice size on C_p distribution, for a NACA 23012 airfoil at an angle of attack of 3° and with a quarter-round ice accretion at 10% chord.

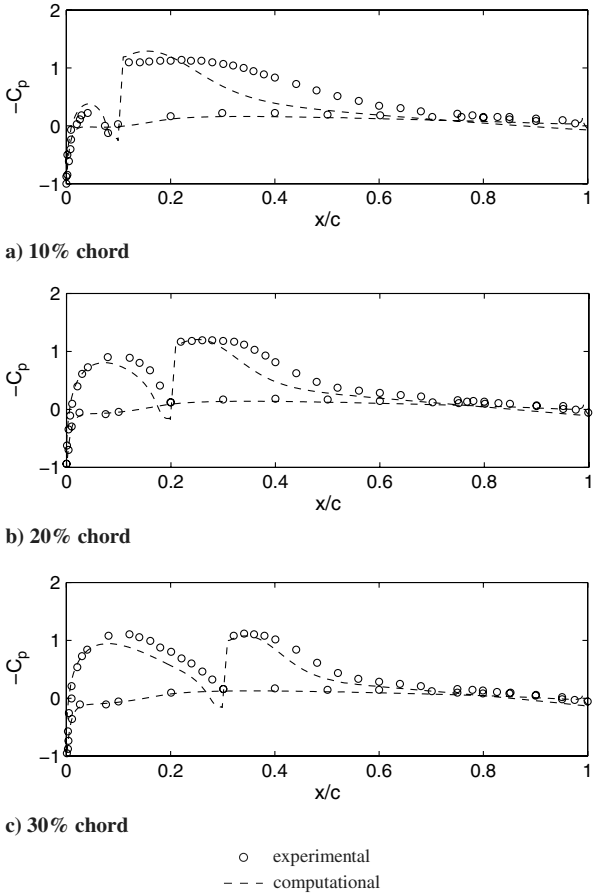

Fig. 10 Effect of ice location on C_p distribution, for a NACA 23012 airfoil at an angle of attack of 3° with a 0.83% chord quarter-round ice accretion.

quarter-round ice accretion on a NACA 23012 airfoil at an angle of attack of 3°. Again, despite the difficulties in predicting the separation bubble (most evident when the ice is at 10% of the chord, as this is the location with the strongest influence on the aerodynamic characteristics of the airfoil [34]) the trends are correctly predicted. While the airfoil with the quarter-round ice shape did not display a suction peak at the leading edge (the reason for the large loss in lift), the suction peak reappears as the ice shape moves backwards, although the peak is not as large as for the clean airfoil.

These results show that OpenFOAM is able to accurately predict the experimental trends for the lift and drag coefficients, as well as the pressure distributions along the airfoil surface under different icing conditions. The largest discrepancies occur for medium angles of attack (in the presence of a large separation bubble that reattaches before the trailing edge) and are due to the numerical approach used (RANS with the Spalart–Allmaras turbulence model). Even though better predictions could probably be obtained with a more sophisticated approach [30], the current level of accuracy is deemed sufficient for the purposes of this work.

IV. Design Optimization System

The capability to correctly predict a physical phenomenon, although fundamental for its understanding, does not in itself improve the performance of a system [10]. Most real-world design problems are characterized by large numbers of design variables, of (often conflicting) objectives and of constraints, which complicate the search for optimal solution(s). In complex design problems, finding the optimal solution by trial and error is extremely difficult if not impossible [35]: an automated search of the design space (design optimization) can facilitate a better exploration of the available options, with a simultaneous reduction in design times. Apart from the numerical optimizer, the essential components of a design optimization system are an efficient parameterization of the design space, which expresses the design task in terms of a number of design variables (forming the design vector) that can be changed by the optimizer, and the evaluation tools, which attribute figures of merit to each candidate design vector. Most real-world design problems involve the evaluation of a number of performance metrics, which often cannot easily be reduced to a single figure of merit; the use of a multi-objective optimizer then allows multiple objectives to be considered simultaneously and leads to identification of a set of equally important design solutions (the Pareto front) for the designer to choose from.

A. Evaluation Tools

Although OpenFOAM has demonstrated the ability to predict with adequate accuracy the aerodynamic performance of an airfoil in both clean and iced conditions, the time required for a single evaluation (30 min on an eight-core Intel Xeon 2.33 GHz processor) in combination with the number of analyses likely to be needed for a satisfactory exploration of the design space make the computational cost of a design optimization prohibitively high. While the use of high-fidelity evaluation tools is essential for a reliable estimation of the candidate designs' figures of merit, and in order not to mislead the optimizer, lower-fidelity analysis tools can be used to undertake a preliminary screening of the design alternatives with the aim of reducing the number of high-fidelity (high-cost) evaluations required. XFoil is a shareware code developed by Drela [36] for the design and analysis of subsonic airfoils, which uses a vorticity-stream function panel method coupled with a two-equation integral boundary layer formulation and a Karman–Tsien correction for compressibility effects. Force coefficients are calculated via the momentum deficit in the wake. Reliable predictions can be obtained

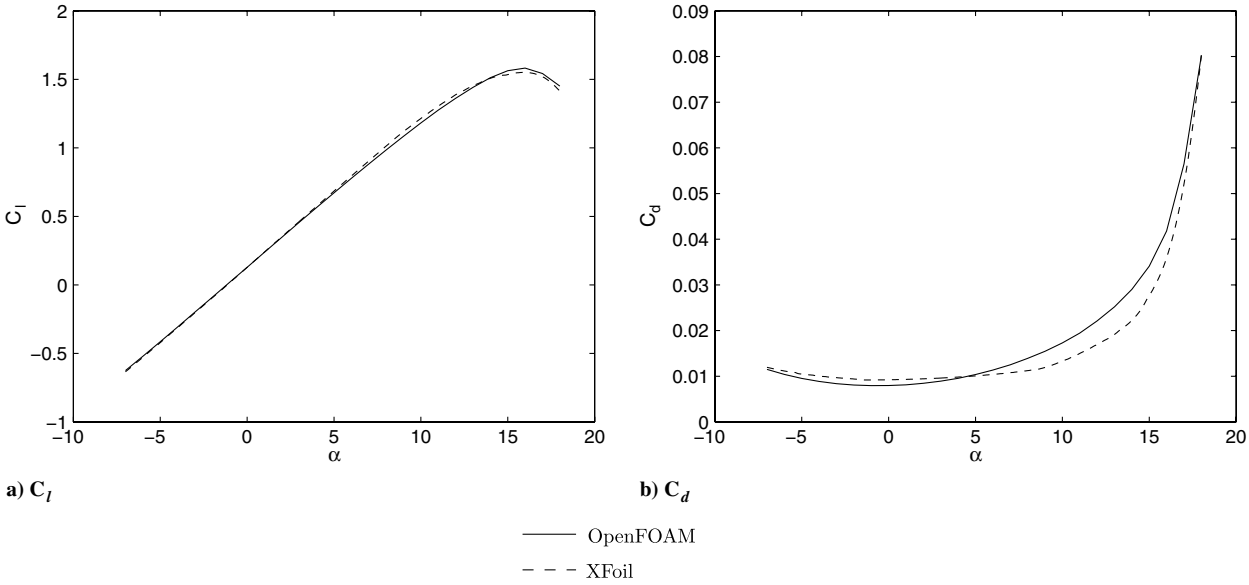

Fig. 11 Comparison of OpenFOAM and XFoil predictions for the clean airfoil.

in the presence of attached flows and even small separated regions, in a fraction of the time required by a Navier–Stokes solver (about a minute for the complete lift curve), making it ideal for the preliminary screening of candidate design.

Figure 11 presents a comparison of the C_l and C_d predictions obtained with OpenFOAM and XFoil. The C_l-α curves are almost indistinguishable, while XFoil underpredicts C_d at high angles of attack.

B. Parameterization Scheme

At the heart of every shape optimization problem is the ability to change the geometry to yield an improvement in the objective functions. The choice of the parameterization scheme is fundamental both in facilitating an effective exploration of the design space and in enabling the desired improvements to be achieved in an acceptable time frame [11]. Given the same evaluation tools, the parameterization approach is what defines the optimization problem formulation and the design space topology [37].

Comprehensive surveys of the different parameterization approaches available for shape optimization are given by Samereh [38], Shahpar [11] and Song and Keane [37]. These include: the use of typical engineering parameters, discrete approaches, basis vectors, domain element methods, polynomials and splines: both Bezier-curves and nonuniform rational Bezier splines, partial differential equations, CAD-based methods, and FFD. FFD is a subset of the soft object animation algorithms used in computer graphics for deforming surfaces and solid models [38]. Mathematically, it is defined in terms of a tensor product trivariate Bernstein polynomial:

$$\mathbf{X}_{\text{FFD}} = \left(\sum_{i=0}^{l}\binom{l}{i}(1-s)^{l-i}s^{i}\right)\left(\sum_{j=0}^{m}\binom{m}{j}\right. \times (1-t)^{m-j}t^{j}\left)\left(\sum_{k=0}^{n}\binom{n}{k}(1-u)^{n-k}u^{k}\right)\mathbf{P}_{ijk}\right. \tag{1}$$

where $\mathbf{X}_{\text{FFD}}$ is the deformed position of an arbitrary point with initial coordinates (s, t, u). This approach assumes a lattice of $l \times m \times n$ initially equally spaced control points. The deformed positions of these nodes are defined by the vectors $\mathbf{P}_{ijk}$ [39].

The main advantages of FFD are its independence from the grid topology (which makes it capable of deforming any type of surface or solid primitive) and its nonintrusiveness (a geometry is represented in terms of the movements of a number of control points from an initial position). FFD is capable of generating radically new shapes, but its

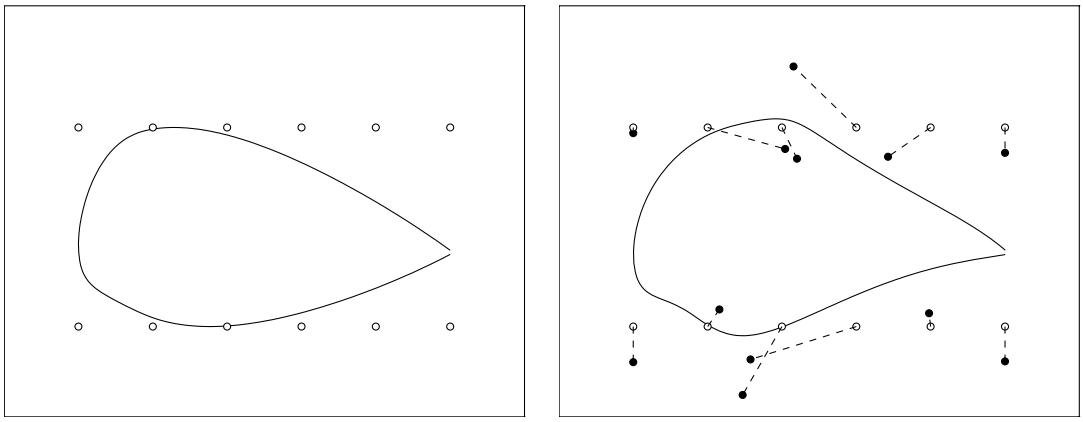

Fig. 12 Example of FFD.

generality can also lead to large design spaces in applications where specialized, highly efficient parameterization schemes can reduce the number of design variables considerably [37].

Figure 12 presents an example of the FFD of a NACA 23012 airfoil. 12 control points ($6 \times 2 \times 1$) were used; those at the extreme left and right were allowed to move only in the y-direction, to keep the airfoil chord fixed, giving a total of 20 design variables.

C. Optimizer

The tabu search (TS) algorithm developed by Jaeggi et al. [40] was selected for this work. TS is a metaheuristic method designed to help a search negotiate difficult regions of the search space (e.g. escape from local minima or cross infeasible regions) by imposing restrictions [41]. The local search phase at its heart is conducted with the Hooke and Jeeves (H&J) algorithm [42], with the best allowed move being made. The short term memory (STM) records the last S visited points, which are tabu and thus cannot be revisited. The effect of the STM is that the algorithm behaves like a normal hill-descending algorithm until it reaches a minimum, then it is forced to climb out of the hollow and explore further. Two other important features of the TS algorithm are intensification and diversification. Intensification is associated with the medium term memory (MTM) where the best M solutions located thus far are stored. Diversification is associated with the long-term memory, which records the areas of the search space that have been explored by dividing it into a number of sectors and tallying how many times each sector has been visited. On diversification the search is restarted in an under-explored region of the design space. The extension to multi-objective problems is straightforward: the MTM contains the set of nondominated solutions found, while at every H&J step a move is selected randomly from among the set of nondominated new designs. The discarded designs are not lost, however: they are stored in the MTM, if appropriate, and can then be selected during intensification.

TS has been found to be particularly effective in the solution of real-world optimization problems. Thanks to the local search algorithm at its heart, it is able to efficiently navigate the complex, highly-constrained search spaces typical of such problems; in contrast, the larger design changes typically made by evolutionary algorithms tend to generate large numbers of infeasible designs [43,44]. This strength can, however, become a weakness of TS in very large design spaces, as H&J search requires a number of evaluations per optimization step approximately equal to twice the number of design variables. Ghisu et al. [45] developed an improvement to the algorithm based on Principal Components' Analysis of the approximation set. This allows the design space parameterization to be rotated towards its most "energetic" directions (those most likely to lead to improved solutions) and the design space dimensionality to be reduced temporarily. By suspending use of the less important (and often more noisy) design variables, this scheme reduces the number of evaluations required at each optimization step and simultaneously improves search effectiveness.

V. Polynomial Chaos

A. Generalized Polynomial Chaos

Real-world problems are often characterized by a number of uncertain parameters. The phenomenon of ice accretion on aircraft surfaces is no exception. A number of factors (such as the water content of the air, droplet size, speed, temperature, angle of attack, etc.) influence the ice formation process (shape, size and position) and hence the aircraft's performance. Neglecting these uncertainties during the design process can lead to a significant difference between anticipated and actual performance [12].

A number of methods exist for assessing the effects of uncertain parameters on the performance of a system. Monte Carlo simulations (MCS), interval analysis, sensitivity derivatives, moment methods and PC are the most common approaches [46]. In recent years, PC has received attention as it offers an exact means (in theory, if an infinite expansion is used) of propagating uncertainties through a system, providing high-order information (similar to MCS) with reduced cost [47].

The original PC formulation, proposed by Wiener [48], employed Hermite polynomials in terms of Gaussian random variables as the basis functions to produce a spectral representation of a random variable. The use of Hermite polynomials is particularly efficient when the input uncertainty is Gaussian-distributed, because it can then be expressed through a first-order hermite chaos (HC) expansion. Lucor et al. [49] demonstrated that an exponential convergence rate can be achieved in this case, while for differently distributed input uncertainties the convergence rate may be substantially slower.

An extension of the HC expansion was proposed by Xiu and Karniadakis [50] to deal with more general random inputs more efficiently. The underlying concept remains the same, but the Hermite polynomials basis is replaced by a generic basis of orthogonal polynomials. A random process $Y(\theta)$ can still be expressed as

$$Y(\theta) = c_0 I_0 + \sum_{i_1=1}^{\infty} c_{i_1} I_1(\xi_{i_1}(\theta)) + \sum_{i_1=1}^{\infty} \sum_{i_2=1}^{i_1} c_{i_1 i_2} I_2(\xi_{i_1}(\theta), \xi_{i_2}(\theta)) + \sum_{i_1=1}^{\infty} \sum_{i_2=1}^{i_1} \sum_{i_3=1}^{i_2} c_{i_1 i_2 i_3} I_3(\xi_{i_1}(\theta), \xi_{i_2}(\theta), \xi_{i_3}(\theta)) + \cdots \tag{2}$$

where $I_n(\xi_{i_1}(\theta), \ldots, \xi_{i_n}(\theta))$ represents the Wiener–Askey polynomial of order n in terms of the random vector $\boldsymbol{\xi} = (\xi_{i_1}, \ldots, \xi_{i_n})$. The polynomials I_n can now be any orthogonal polynomials.

For notational convenience, Eq. (2) can be rewritten as

$$Y(\theta) = \sum_{j=0}^{\infty} \hat{c}_j \Phi_j(\boldsymbol{\xi}) \tag{3}$$

where there is a one-to-one correspondence between the polynomials $I_n(\xi_{i_1}, \ldots, \xi_{i_n})$ and $\Phi_j(\boldsymbol{\xi})$ and between the coefficients $\hat{c}_j$ and $c_{i_1,\ldots,i_n}$.

The PC basis and supporting distribution most appropriate to the given input uncertainty distribution can be chosen from the Askey scheme (see Table 1): this allows (as in the case of Gaussian distributions and HC) certain input uncertainties to be expressed exactly through a first-order expansion, thereby simplifying the task of solving the stochastic differential equations governing the problem, and improving the accuracy of the solution. With this approach, Xiu and Karniadakis [50] achieved exponential convergence for a number of different input distributions.

1. *Nonintrusive Polynomial Chaos*

In complex applications, a direct modification of the mathematical system of equations describing the problem is often far from straightforward. Le Maître et al. [51] developed a NIPC alternative, in which the different modes of the expansion can be reconstructed by projecting deterministic computations onto the PC basis:

$$\hat{c}_i = \frac{1}{\langle \Phi_i^2 \rangle} \int \mathrm{d}\xi_1 \cdots \int \mathrm{d}\xi_N Y(\boldsymbol{\xi}) \Phi_i(\boldsymbol{\xi}) W(\boldsymbol{\xi}) \tag{4}$$

Table 1 Orthogonal polynomials from the Askey scheme [50]

Random variable	PC basis
Continuous	
Gaussian	Hermite
Gamma	Laguerre
Beta	Jacobi
Uniform	Legendre
Discrete	
Poisson	Charlier
Binomial	Krawtchouk
Negative binomial	Meixner
Hypergeometric	Hahn

where $W(\boldsymbol{\xi})$ is the weighting function corresponding to the PC basis Φ_i. As $Y(\boldsymbol{\xi})$ and $\Phi_i(\boldsymbol{\xi})$ are both polynomials, the preceding integral can be calculated exactly through a Gaussian quadrature formula:

$$\hat{c}_i = \frac{1}{\langle \Phi_i^2 \rangle} \sum_{n_1=1}^{N_0} \cdots \sum_{n_N=1}^{N_0} Y(\xi_{n_1}, \cdots, \xi_{n_N}) \Phi_i(\xi_{n_1}, \cdots, \xi_{n_N}) \prod_{k=1}^{N} w_{n_k} \quad (5)$$

where the couples (ξ_k, w_k) represent the one-dimensional Gaussian quadrature points and weights. This formula is exact when the integrand is a polynomial of degree less than or equal to $2N_0 - 1$ [52]. The quadrature points are the zeros of the orthogonal polynomial of degree N_0 and the weights can be calculated as

$$w_k = \frac{1}{[\phi'_{N_0}(\xi_k)]^2} \int_{-\infty}^{+\infty} W(\xi) \left[\frac{\phi_{N_0}(\xi)}{\xi - \xi_k} \right]^2 \mathrm{d}\xi \quad (6)$$

in which $\phi'_n(\xi_k)$ is the first-order derivative of the orthogonal polynomial of order n at quadrature point ξ_k.

It is important to note that the number of deterministic calculations required by the NIPC approach is equal to $(p+1)^N$, where N is the number of random dimensions and p the chaos order. This is always greater than the $\frac{(p+N)!}{p!N!}$ calculations required by intrusive PC [51], but the possibility of treating the analysis software as a black box is a significant advantage.

B. Dealing with Nonstandard Distributions

The Cameron–Martin theorem [53] ensures the possibility of using HC to represent any second-order process. In principle, any set of orthogonal polynomials can be used to express any distribution. Choosing the weighting function to match the random variable's input distribution (and choosing the PC basis as a consequence) has the advantage that the input random variable can be expressed exactly as a first-order expansion with respect to the PC basis, improving convergence and simplifying the computations.

In the case of a nonstandard distribution the input variable needs to be expressed as a function of the support variable (or, equivalently, the coefficients of its expansion in terms of the elements of the PC basis need to be determined). As for NIPC, the coefficients can be obtained by projecting the generic input variable Y onto the PC basis:

$$\hat{c}_i = \frac{\langle Y \Phi_i \rangle}{\langle \Phi_i^2 \rangle} = \frac{1}{\langle \Phi_i^2 \rangle} \int Y \Phi_i(\xi) W(\xi) \, \mathrm{d}\xi \quad (7)$$

This approach works reasonably well for standard random variables [50], but it becomes problematic for nonstandard ones (e.g. variables with multimodal distributions [13]), requiring a high-order expansion that leads to a large number of coupled equations in the intrusive formulation or high-order quadratures in the nonintrusive method.

A more effective approach for dealing with generic random distributions is to numerically construct the basis of orthogonal polynomials with respect to the given random distributions; a detailed review of the main techniques available for generating orthogonal polynomials (based on the Stieltjes procedure or the Chebyshev algorithm) is given by Gautschi [54]. Wan and Karniadakis [55] used this approach to decompose a standard random input into a number of smaller elements by defining a conditional probability function, thereby significantly improving the convergence characteristics of PC in the presence of long-term integrations. Ghisu et al. [13] extended the method to deal with generic polynomial input distributions, achieving good convergence characteristics in the presence of both one-dimensional and multidimensional uncertainties.

VI. Results

A. Initial Optimization: Lift and Drag

The objective of this work is to design airfoils that are less sensitive to ice accretion. However, if possible, this should not be achieved at the cost of a degradation in the airfoil's performance when clean, as this represents the condition under which an aircraft flies most of the time. If this is not possible, it is important that the tradeoff between clean and iced performance is clearly identified. Considering the relatively low probability of ice accretion but also recognising the risks linked to a modification of the airfoil's effective shape, clean and iced airfoil performance represent incommensurable design objectives, which are therefore effectively handled by means of a multi-objective optimization approach. In this study clean C_d is used as a measure of the airfoil's clean performance, while the degradation in C_l in the presence of a 1% chord quarter-round ice shape located at 15% chord is used to measure the impact of ice accretion. This ice shape was selected because of its influence on an airfoil's aerodynamic performance (relative to other tested shapes [26]) and its size as representative of large ice formations [34]. Thus the accretion represents the most extreme shape and size likely to be encountered. The 15% chord location was selected taking into consideration both the location causing maximum performance degradation on the NACA 23012 airfoil (10% chord) and the extent of the region where ice deposition is most likely to occur (in the front 40% of the airfoil)[34]. Inlet Mach and Reynolds numbers were kept fixed, respectively, at 0.18 and 1.8×10^6 as in the validation studies.

As evaluation of the entire lift curve is impractical due to time constraints, a single condition (corresponding to a clean C_l of unity) is considered. XFoil is used to evaluate the clean airfoil's performance (and the angle of attack for the iced performance simulations), while OpenFOAM is used to evaluate the impact of ice accretion. The maximum C_l of the clean airfoil is constrained not to be lower than that of the datum airfoil (a NACA 23012). Constraints are also set for the airfoil's 25 and 50% thickness-to-chord, its leading-edge radius and trailing-edge angle (for structural reasons). The optimization problem is summarized in Eq. (8), in which an overbar indicates a value associated with the datum airfoil:

$$\begin{gathered} \text{Minimize} \quad C_{d,\text{clean}} \quad \Delta C_l \\ \text{Subject to} \quad C_{l,\text{clean,max}} \geq \overline{C}_{l,\text{clean,max}} \\ (t/c)_{25\%} \geq \overline{(t/c)}_{25\%} \quad (t/c)_{50\%} \geq \overline{(t/c)}_{50\%} \\ r_{\text{LE}} \geq \overline{r}_{\text{LE}} \quad \beta_{\text{TE}} \geq \overline{\beta}_{\text{TE}} \end{gathered} \quad (8)$$

Design solutions which violate one or more constraints, as well as those resulting in nonconverged simulations, are treated with a barrier function (i.e. a high value is assigned to the figures of merit), thus suggesting that the optimizer move away from the corresponding design.

Three hundred optimization steps were completed, for a total of 7,071 evaluations. Sixty-two percent of these were analyzed with OpenFoam, as they passed the preliminary screening (geometrical feasibility and low-fidelity analysis with XFoil), requiring 11 days on an eight-node cluster (each machine being an eight-core Intel Xeon 2.33 GHz processor). The results for this first optimization problem

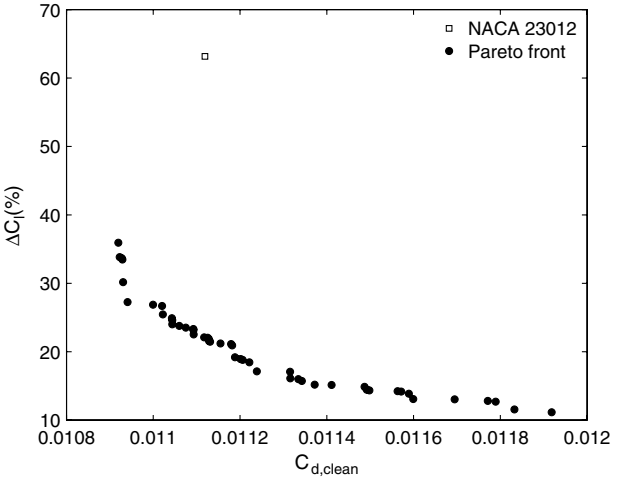

Fig. 13 Pareto front for the two-objective optimization problem [Eq. (8)].

are presented in Fig. 13 and show that significant, simultaneous improvements in both clean and iced airfoil performance are possible. Figure 14 compares the datum airfoil with the two designs (for minimum clean C_d and minimum ΔC_l) at the ends of the Pareto front. Figure 15 compares the C_p distributions for these three airfoils, respectively, in clean and iced conditions.

The best C_d configuration achieves a 1.9% reduction in drag thanks to a less pronounced pressure peak and adverse pressure gradient. These allow a less pronounced boundary layer development and hence a reduction in the profile drag, while the fraction of drag due to skin friction is virtually unchanged. The presence of a lower suction peak and of a region of zero pressure gradient just after 15% chord reduces the effect of an ice accretion in this region significantly, as shown in Fig. 15b. This simultaneously allows a lower pressure increase in the stagnation region ahead of the ice and a larger pressure peak downstream, which together result in a larger lift coefficient. The situation is even more marked for the best iced C_l (or lower ΔC_l) design configuration (Fig. 15c), which features an indentation at about 15% chord. This produces a more pronounced adverse pressure gradient just after the leading edge, followed by a negative pressure gradient which helps the flow to reattach after the accretion, improving the iced performance of the airfoil. This is obtained at the price of a 7% increase in the overall drag of the airfoil, mainly caused by a deterioration in its profile drag. While this deterioration in clean airfoil performance would probably be unacceptable in practice, C_d for the clean airfoil was not constrained during this optimization in order not to reduce further the size of the feasible region of the design space. This allows the optimizer to identify a variety of airfoil shapes, which could contain, at least in part, some of the features that can contribute to improved overall performance. Forty-four equal quality (from a multi-objective perspective) designs can be selected from the Pareto front in Fig. 13: 17 of these dominate the datum design (i.e. they are better with respect to both objectives), while the remaining 31 achieve their improvements in ΔC_l at the price of deteriorations in clean airfoil performance.

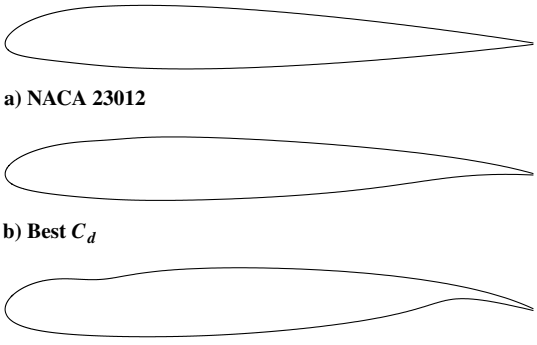

a) NACA 23012

b) Best C_d

c) Best ΔC_l

Fig. 14 Comparison of datum airfoil with the extreme Pareto-optimal designs.

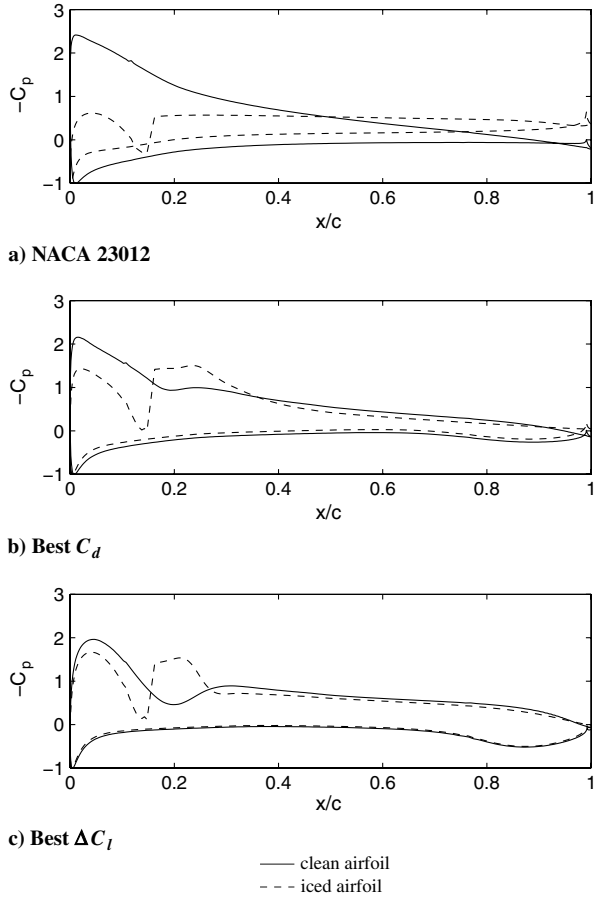

a) NACA 23012

b) Best C_d

c) Best ΔC_l

Fig. 15 Comparison of clean and iced (with a 1% chord quarter-round ice accretion at 15% chord) pressure coefficient distributions, for the datum airfoil and the two extreme designs.

B. Second Optimization: Lift, Drag, and Pitching Moment

It is evident from Fig. 15 that the improvement in iced airfoil performance was obtained by reducing its forward loading (i.e. by decreasing the amount of lift produced in the leading-edge region) at the price of increased negative pitching moments. The consequence of too large a negative (or nose-down) pitching moment is that the horizontal tail has to provide the down force necessary to balance the aircraft, resulting in a large trim drag. The existence of a large negative pitching moment affects the location of the aircraft's center of gravity, which may need to be positioned further aft to ensure sufficient static stability. It is therefore important to use airfoil sections with a low negative pitching moment, especially on small aircraft, which have fewer means of counteracting a nose-down tendency. It is clear that iced C_l, clean C_d and pitching moment coefficient C_m represent competing objectives, which can be treated effectively by means of a three-objective optimization. The resulting optimization problem is summarized in Eq. (9); again an overbar indicates a value associated with the datum airfoil:

$$\begin{aligned} &\text{Minimize} \quad C_{d,\text{clean}} \quad \Delta C_l \quad -C_{m,\text{clean}} \\ &\text{Subject to} \quad C_{l,\text{clean,max}} \geq \overline{C}_{l,\text{clean,max}} \\ &(t/c)_{25\%} \geq \overline{(t/c)}_{25\%} \quad (t/c)_{50\%} \geq \overline{(t/c)}_{50\%} \\ &r_{\text{LE}} \geq \overline{r}_{\text{LE}} \quad \beta_{\text{TE}} \geq \overline{\beta}_{\text{TE}} \end{aligned} \tag{9}$$

Five hundred optimization steps were performed, for a total of 11,327 configurations (59% of which analyzed with OpenFoam in 17 days). Figure 16 presents the results from the three-objective optimization: the black dots represent the nondominated (Pareto-optimal) designs, while the light gray dots show their projections onto the possible two-objective planes. The datum design (NACA 23012) is represented with a gray diamond, while the projections on the two-objective planes are shown with black dashed lines. Figure 17 presents the same tradeoff surface in a two-dimensional plot (C_d vs ΔC_l), with the gray scale indicating the $-C_m$ value. The position of the datum design is again represented by the two black dashed lines. It is evident that considerable improvements in iced airfoil performance can be achieved without deteriorating clean performance and with only a limited impact on the airfoil's pitching moment.

C. Robust Optimization

In the previous subsection, the performance of an airfoil subject to a quarter-round ice accretion at 15% chord (representing a ridge accretion in SLD or runback ice conditions) was optimized.

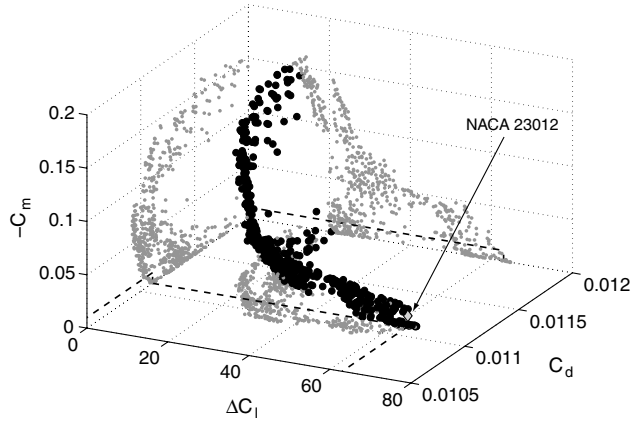

Fig. 16 Pareto front for the three-objective optimization problem [Eq. (9)].

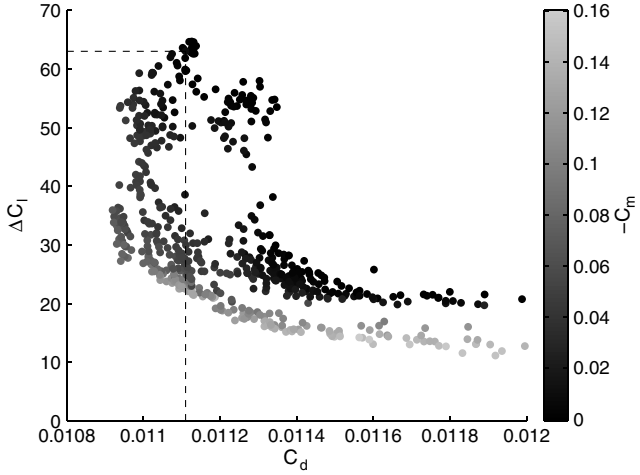

Fig. 17 Two-dimensional representation of Pareto front for the three-objective optimization problem [Eq. (9)].

However, as for any real-world problem, the formation of ice on any of the aircraft's surfaces is influenced by a number of uncertain parameters which determine its size, position, shape, roughness, etc.

As previously mentioned, the shape and size of the ice accretion were chosen to be representative of the most extreme conditions likely to be encountered. Of the remaining uncertain parameters, the ice location is the one with the largest influence on airfoil performance, while the effects of other parameters (such as surface roughness, Reynolds and Mach numbers) are comparatively small [33]. In the optimizations presented thus far, the location of the ice formation has been kept fixed at a nominal value, chosen because of its large influence on the performance of a NACA 23012 airfoil; other airfoils are more sensitive to the formation of ice at different positions [34]). A design approach that considers only nominal values of uncertain parameters risks producing an extreme design: one with improved performance for the given values of the uncertain parameters, but which is very sensitive to any variation in these parameters. This is evident from the solutions in Fig. 14, which achieve optimal performance in the presence of a quarter-round ice shape at 15% chord, but are sensitive to the formation of ice closer to the leading edge, where there is a strong adverse pressure gradient. For instance, of the 44 solutions in Fig. 13 and 18 will not yield a converged flow solution for ice at 5% chord, implying poor performance. The remaining 26 have a worse value of ΔC_l for ice at 5% chord than at 15% (with the change in ΔC_l being as much as 72%).

To find airfoil shapes less sensitive to the location of the ice accretion, a further optimization problem was defined, in which the position of the quarter-round ice shape was considered uncertain, with a given PDF. The PDF shown in Fig. 18 was chosen to cover the area where ice accretion is most likely to occur, but does not reflect any real ice formation data. It is used here to demonstrate the generality of the method and could be replaced by any PDF more closely representing actual ice formation data, without requiring any modifications to the method. The optimization was run for 500 further steps, for a total of 11,511 evaluations (60.5% with OpenFoam). The total run-time was six weeks on the eight-node eight-core Intel Xeon cluster.

A NIPC formulation was used to calculate the variability in airfoil performance in the face of the given variation in accretion position. This method (introduced in Sec. V.A) offers a convenient means of propagating uncertainties through an otherwise deterministic system, providing high-order information (similar to MCS) at a much reduced cost. A previous study by the authors [13] shows how the use of an adaptive formulation (where the orthonormal polynomials defining the chaos basis and the quadrature rules are calculated in real time relative to the given PDF) facilitates improved convergence (and hence a reduction of the number of deterministic simulations required).

The support variable ξ was chosen with the same distribution as L, normalized between 0 and 1 to simplify the derivations that follow. This means that the input uncertainty can be expressed exactly through a first-order PC expansion. The coefficients of the PC

Fig. 19 NACA 23012 with ice accretions at the quadrature nodes.

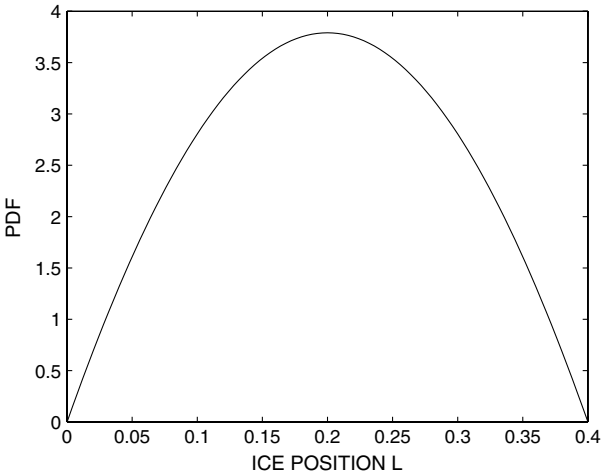

Fig. 18 PDF for the ice location *L*.

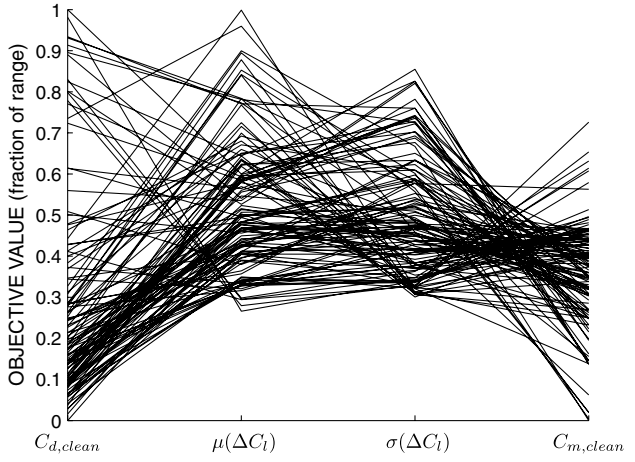

Fig. 20 Parallel coordinates representation of the feasible designs from the second optimization.

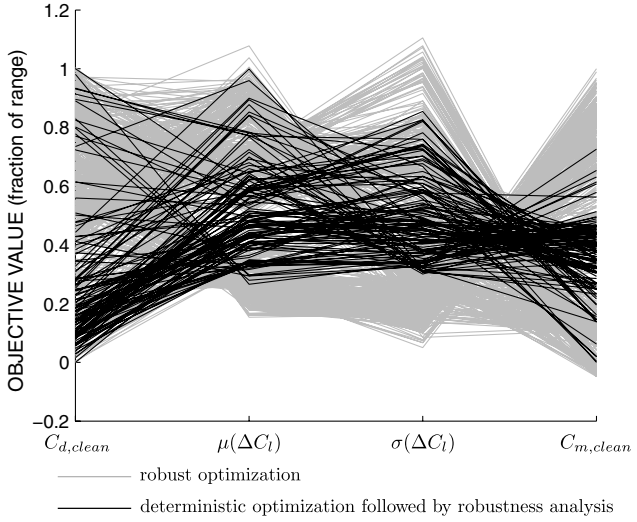

Fig. 21 Pareto front resulting from the robust optimization problem in Eq. (12).

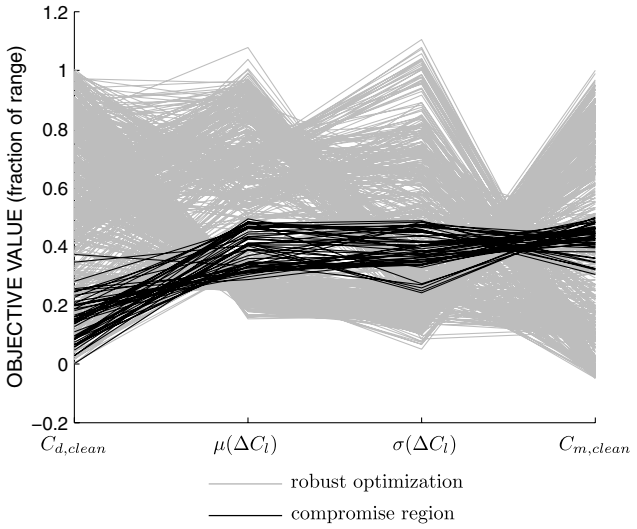

Fig. 22 The compromise region of the Pareto front.

Table 2 System of orthonormal polynomials

Order	P_i^*	$\langle (P_i^*)^2 \rangle$
0	1	1
1	$2.2500 - 4.5000\xi$	1
2	$3.7808 - 18.8458\xi + 18.8458\xi^2$	1
3	$5.5484 - 49.7343\xi + 115.9126\xi^2 - 77.2751\xi^3$	1
4	$7.5249 - 104.8989\xi + 418.8507\xi^2$ $-627.9036\xi^3 + 313.9518\xi^4$	1

Table 3 Roots and weights of a fourth-order quadrature

Root ξ_i	Weight w_i	P_0^*	P_1^*	P_2^*	P_3^*	P_4^*
0.1183	0.1164	1.0000	1.7178	1.8155	1.1599	0.0
0.3580	0.3836	1.0000	0.6389	−0.5508	−0.9462	0.0
0.6420	0.3836	1.0000	−0.6389	−0.5508	−0.9462	0.0
0.8817	0.1164	1.0000	−1.7178	1.8155	−1.1599	0.0

expansion for any dependent variable can be calculated as in Eq. (5). This requires the determination of the basis of orthogonal (or orthonormal) polynomials and of the quadrature rule (nodes and weights). The first five orthonormal polynomials are shown in Table 2, while the quadrature nodes and weights of a fourth-order quadrature are shown in Table 3.

Ice accretions at the quadrature nodes are shown in Fig. 19, for a NACA 23012 airfoil.

Given the coefficients of the PC expansion, the relative moments can readily be evaluated and used as objectives in an optimization problem. Mean and variance can be calculated as follows:

$$\mu = \int \cdots \int \sum_{i=0}^{p} y_i P_i^*(\xi) W(\xi)\, \mathrm{d}\xi = y_0 \tag{10}$$

$$\sigma^2 = \int \cdots \int \left(\sum_{i=0}^{p} y_i P_i^*(\xi) - y_0 \right)^2 W(\xi)\, \mathrm{d}\xi = \sum_{i=1}^{p} y_i^2 \tag{11}$$

To define a suitable robust optimization problem, the deterministic value of ΔC_l in Eq. (8) was replaced by two objectives (the mean and standard deviation of ΔC_l). The resulting four-objective optimization problem is summarized in Eq. (12):

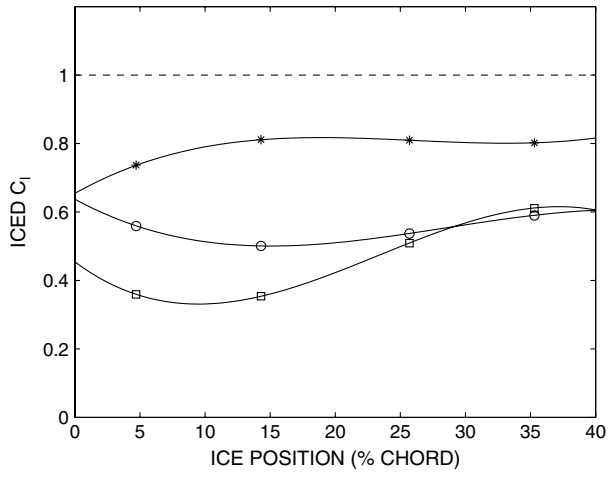

a) Dependence of C_l on ice location

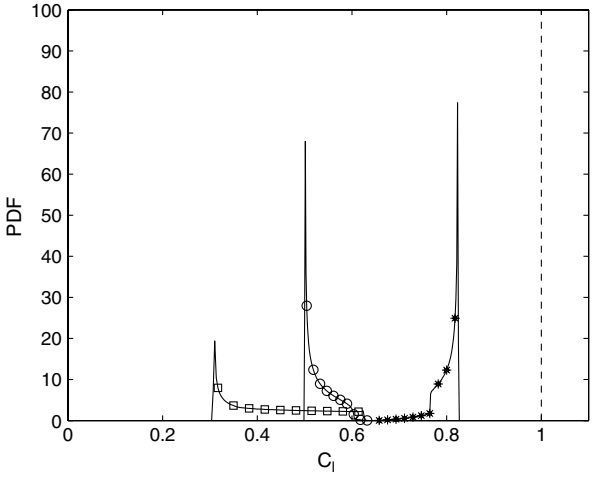

b) C_l probability density functions

clean C_l
NACA 23012
compromise design
minimum mean ΔC_l design

Fig. 23 Iced C_l variation as a function of the ice position.

$$\begin{aligned} \text{Minimize} \quad & C_{d,\text{clean}} \quad \mu(\Delta C_l) \quad \sigma(\Delta C_l) \quad -C_{m,\text{clean}} \\ \text{Subject to} \quad & C_{l,\text{clean,max}} \geq \overline{C}_{l,\text{clean,max}} \\ & (t/c)_{25\%} \geq \overline{(t/c)}_{25\%} \quad (t/c)_{50\%} \geq \overline{(t/c)}_{50\%} \\ & r_{\text{LE}} \geq \overline{r}_{\text{LE}} \quad \beta_{\text{TE}} \geq \overline{\beta}_{\text{TE}} \end{aligned} \tag{12}$$

First, the results from the three-objective deterministic optimization (Fig. 16) were analyzed using the NIPC approach: only 313 designs (48.6% of the total) were found to be feasible (i.e. a converged solution could be obtained for each of the quadrature node ice positions shown in Fig. 19. The objective values for the nondominated of these (125 designs) are shown in Fig. 20, in a parallel coordinates representation: each line connects the objective values of a particular design. The first and fourth objective values have been normalized using the minimum and maximum values found during the optimization. The minimum and maximum values of ΔC_l have been used to normalize $\mu(\Delta C_l)$, while $\sigma(\Delta C_l)$ was normalized using an arbitrarily chosen value of $\Delta C_l/4$.

This postoptimization analysis of robustness allows the identification of designs with robust behavior over the range of variability of the uncertain variable and the discarding of designs that do not meet this requirement.

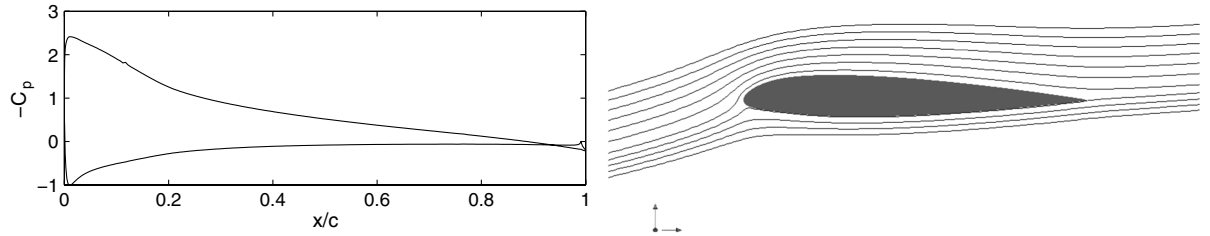

a) Clean airfoil

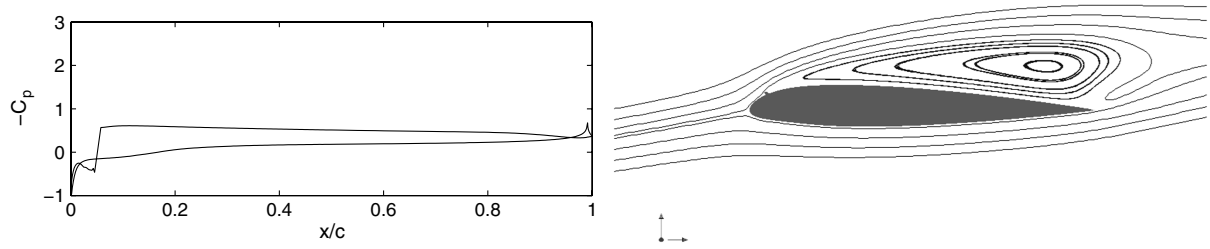

b) 1st quadrature point

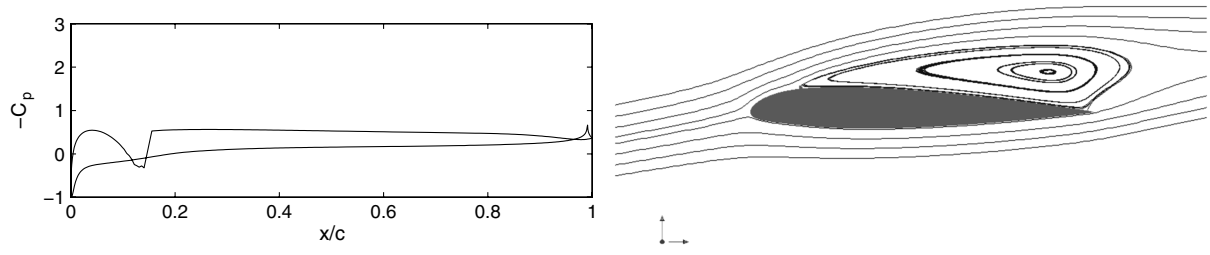

c) 2nd quadrature point

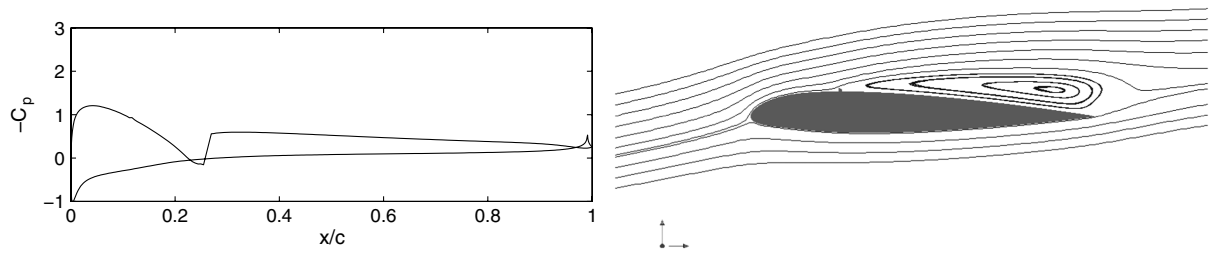

d) 3rd quadrature point

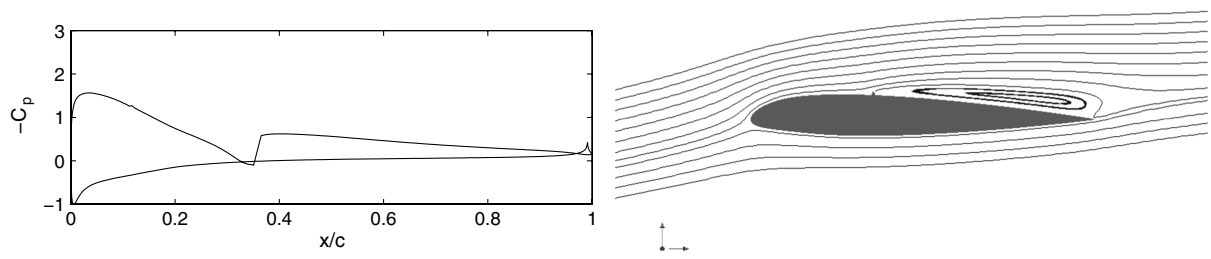

e) 4th quadrature point

Fig. 24 Pressure coefficient distribution and streamlines for the NACA 23012 airfoil, clean and with a quarter-round ice accretion at the quadrature nodes.

An optimization was then initiated using these designs as starting points, searching for further improvements in the new figures of merit: the feasible designs were added to the MTM (which stores the nondominated solutions) and the optimization restarted from one of these configurations, randomly selected. Figure 21 compares the results from this robust optimization (again normalized with the maximum and minimum objective function values found in the three-objective optimization) with the results from the deterministic optimization (after robustness analysis).

Figure 22 highlights the designs with values of every objective function in the lower 50% of the respective range (normalized with the maximum and minimum objective function values found in the three-objective optimization). These represent the compromise area of the Pareto front, which contains designs with good performance for every objective, i.e., the designs most likely to be selected by the decision-maker.

Figure 23 compares the lift coefficients of a compromise design and of the minimum mean ΔC_l design with that of the NACA 23012 airfoil, in the presence of the quarter-round ice accretion. Figure 23a presents the dependence of C_l on ice location, while Fig. 23b shows the C_l PDFs for the three airfoils. The improved iced performance of the best C_l configuration (with a mean ΔC_l of 0.2) is obtained at the

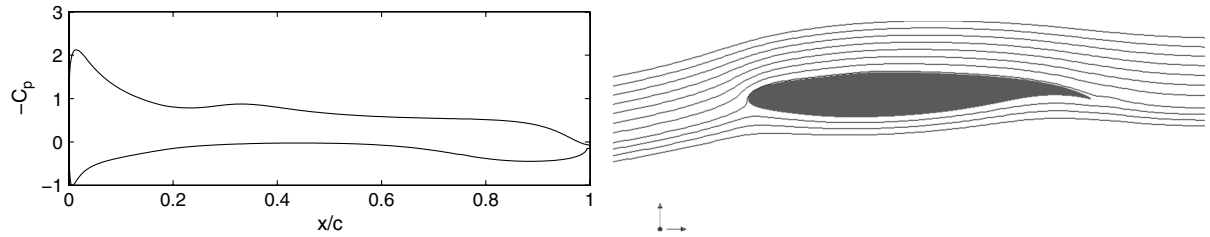

a) Clean airfoil

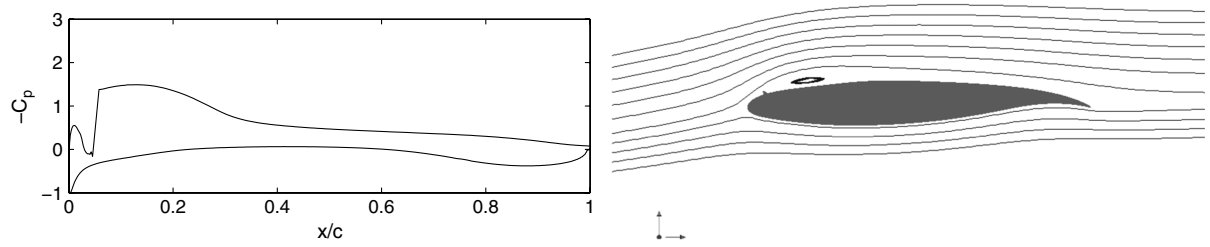

b) 1st quadrature point

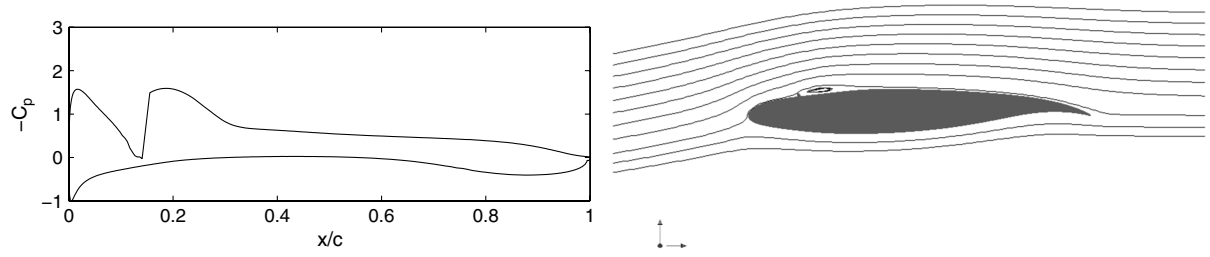

c) 2nd quadrature point

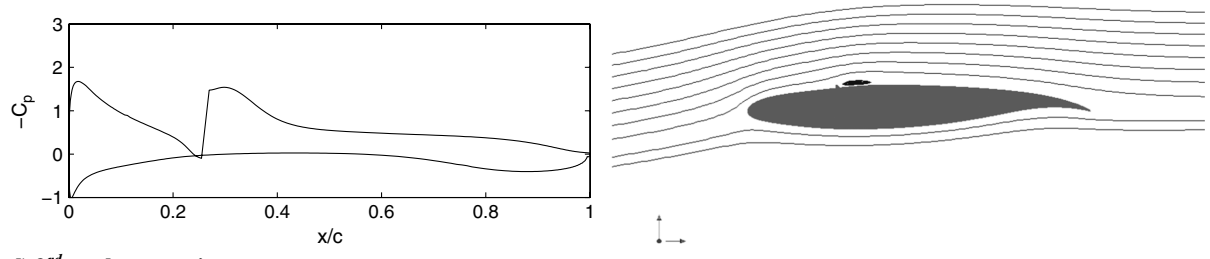

d) 3rd quadrature point

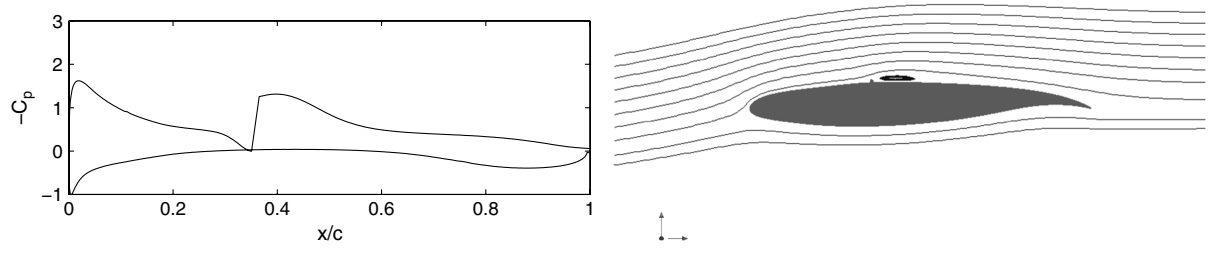

e) 4th quadrature point

Fig. 25 Pressure coefficient distribution and streamlines for the minimum mean ΔC_l airfoil, clean and with a quarter-round ice accretion at the quadrature nodes.

price of a 10% increase in clean C_d (mainly due to a larger profile drag) and a large negative pitching moment. The compromise design achieves a smaller improvement in iced performance (the mean ΔC_l is 0.47 and still compares favorably with the value of 0.57 for the NACA 23012 airfoil) but a −2.1% lower clean C_d than that of the datum design, with a contained negative pitching moment (60% lower than that of the best C_l design). Crucially, the standard deviation for ΔC_l of the compromise design (0.03) is significantly lower than that of the NACA 23012 airfoil (0.10) thanks to less pronounced performance degradation for an ice accretion in the first 20% of the chord. The minimum mean ΔC_l design has a similar standard deviation in ΔC_l to the compromise design (0.029), but a much smaller mean ΔC_l value, of course.

Figures 24–26 present a comparison of pressure coefficient distributions and streamlines for the three airfoils, for different locations of the quarter-round ice shape (corresponding to the quadrature nodes) and an angle of attack giving a clean C_l of 1. The effect of the ice accretion on the NACA 23012 is the reduction (or elimination for ice accretions close to the leading edge) of the leading-edge suction peak present with the clean airfoil (see Fig. 24a), resulting in a drastic reduction in lift. The flow is largely separated (especially for the first two ice accretion locations).

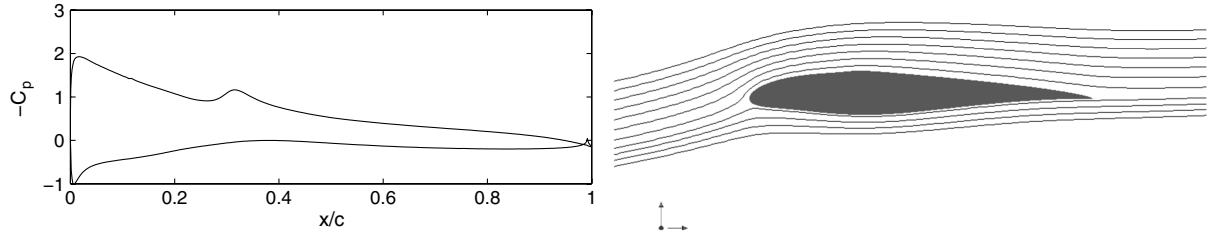

a) Clean airfoil

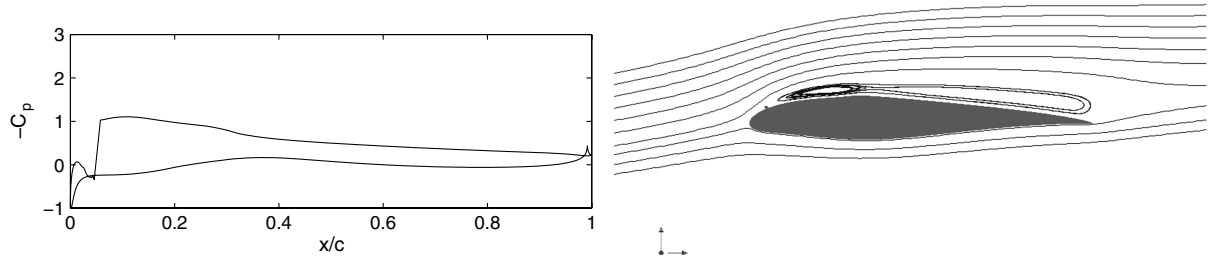

b) 1st quadrature point

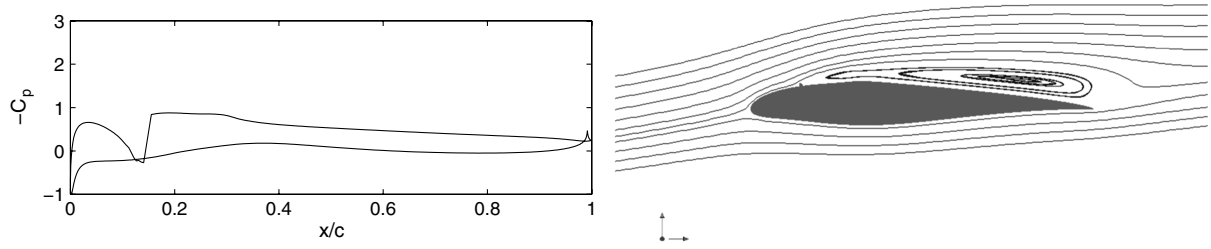

c) 2nd quadrature point

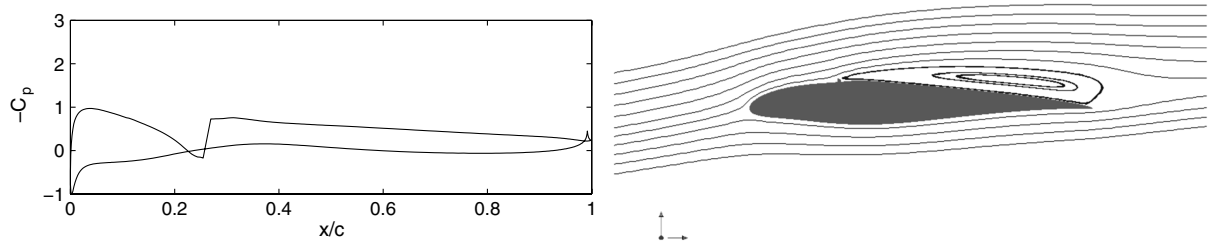

d) 3rd quadrature point

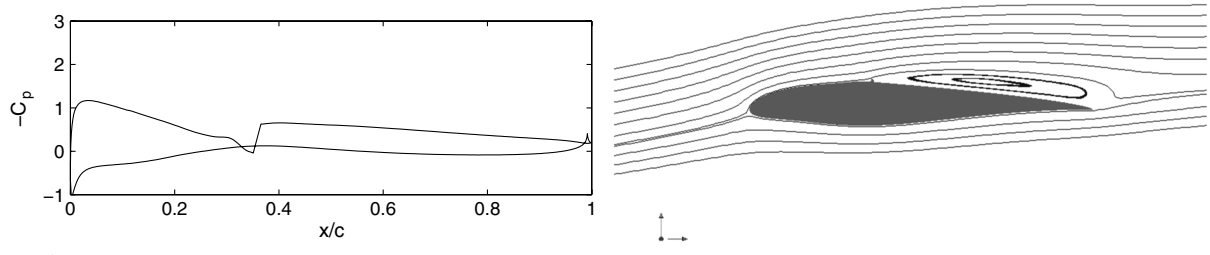

e) 4th quadrature point

Fig. 26 Pressure coefficient distribution and streamlines for the compromise airfoil, clean and with a quarter-round ice accretion at the quadrature nodes.

The minimum mean ΔC_l airfoil (Fig. 25) has a clean pressure distribution with a less pronounced peak and lower adverse pressure gradient (but an increased rear loading which generates the large negative pitching moment), and with a plateau in the suction surface's pressure coefficient between 20 and 35% of the chord. The milder adverse pressure gradient reduces the impact of the ice accretion: the flow rapidly reattaches after the accretion, modifying only marginally the clean pressure coefficient distribution.

The compromise airfoil (Fig. 26) also has a lower suction peak compared with the NACA 23012. A reduced adverse pressure gradient, combined with a region of favorable pressure gradient, help to reduce the impact of the presence of ice accretions. The differences in the size of the separations are evident for the first two quadrature points (those corresponding to the worst NACA 23012 performance), while the performance of the two airfoils for ice at the remaining quadrature nodes are comparable.

It is important to note that the two airfoils considered in detail here represent only two out of 1006 Pareto-equivalent design solutions found in the robust optimization. Depending on the specific application (type of aircraft, available de- or anti-icing systems), the designer can select the most suitable configuration, trading off clean performance (C_d and C_m) against sensitivity to ice ($\mu(\Delta C_l)$ and $\sigma(\Delta C_l)$).

VII. Conclusions

Ice can form when droplets of water below freezing temperature (supercooled) hit an aircraft's surface, freezing upon impact. Ice accretions change the effective shape of aerodynamic surfaces, causing a modification of the forces and moments controlling the aircraft's flight. The dangers of flying into icing conditions have been well recognized since the 1930s and significant progress has been made to reduce the likelihood of ice formation and the associated risks through the implementation of anti-icing and de-icing systems, strict certification procedures and specific pilot training. These advances have been made possible by continuous efforts to understand both the physics of ice accretion and its effects on aircraft aerodynamics, both through experimental and numerical studies. Yet incidents and accidents induced by ice formation continue to occur on all classes of aircraft.

To take full advantage of improved predictive capabilities, the associated analysis methods must be integrated within the design process. In this work, a CFD analysis of an airfoil in the presence of a quarter-round ice shape (simulating an ice accretion in SLD or runback conditions, the most dangerous in terms of performance degradation) has been integrated within an optimization system featuring a metaheuristic multi-objective optimizer and a geometry modeler based on a FFD of an initial airfoil. In recognition of the real-world variability in the ice accretion location (and of its strong influence on aerodynamic performance), this parameter has been treated as uncertain, with a given PDF. An adaptive nonintrusive PC formulation (based on a set of orthonormal polynomials constructed in real time relative to the input PDF) has been used to calculate the consequent variability in the airfoil's performance.

The multi-objective optimization approach adopted allows the identification of a family of solutions (representing the tradeoffs between the conflicting objectives) for the decision-maker to choose from. The results presented demonstrate how significant and robust improvements in iced airfoil performance can be obtained without excessively penalizing its clean performance. In this study, only the ice location has been considered uncertain because of its large impact on airfoil performance, compared with other uncertain parameters such as ice shape size and form, Reynolds and Mach number; additional uncertain parameters can be included without requiring modifications to the analysis and optimization system [13], but with an impact on the computational cost of the optimization process.

References

[1] Johnson, C. L., "Wing Loading, Icing and Associated Aspects of Modern Transport Design," *Journal of the Aeronautical Sciences*, Vol. 8, Dec. 1940.

[2] "Search for Answers Begins in Buffalo Plane Crash, CNN website, http://www.cnn.com/2009/US/02/13/plane.crash.new.york/index.html [retrieved 13 Feb. 2009].

[3] Ratvasky, T. P., Barnhart, B. P., and Lee, S., "Current Methods Modeling and Simulating Icing Effects on Aircraft Performance, Stability, Control," *Journal of Aircraft*, Vol. 47, No. 1, 2010, pp. 201–211.

[4] *Aircraft Icing Handbook*, Civil Aviation Authority, Lower Hutt, New Zealand, 2000.

[5] Aircraft Owners and Pilots Association Air Safety Foundation, *Safety Advisor-Aircraft Icing*, AOPA, Frederick, MD, 2008.

[6] Dunn, T. A., Loth, E., and Bragg, M. B., "Computational Investigation of Simulated Large-Droplet Ice Shapes on Airfoil Aerodynamics," *Journal of Aircraft*, Vol. 36, No. 5, 1999, pp. 836–843. doi:10.2514/2.2517

[7] Bragg, M. B., "Aerodynamics of Supercooled-Large-Droplet Ice Accretions and the Effect on Aircraft Control," *Proceedings of the FAA International Conference on Aircraft Inflight Icing*, Vol. 2, 1996, pp. 323–362; also U.S. Dept. of Transportation/Federal Aviation Administration AR-96/81.

[8] Marongiu, C., Vitagliano, P. L., Zanazzi, G., and Narducci, R., "Aerodynamic Analysis of an Iced Airfoil at Medium/High Reynolds Number," *AIAA Journal*, Vol. 46, No. 10, 2008, pp. 2469–2478. doi:10.2514/1.34550

[9] Wright, W. B., Potapczuk, M. G., and Levinson, L. H., "Comparison of LEWICE and GlennICE in the SLD Regime," AIAA Paper 2008-0439, 46th Aerospace Sciences Meeting, Reno, NV, 2008.

[10] Denton, J. D., and Dawes, W. N., "Computational Fluid Dynamics for Turbomachinery Design," *Journal of Mechanical Engineering Science*, Vol. 213, No. 2, 1999, pp. 107–124. doi:10.1243/0954406991522211

[11] Shahpar, S., "Automatic Aerodynamic Design Optimisation of Turbomachinery Components: An Industrial Perspective," *VKI Lecture Series on Optimisation Methods and Tools for Multicriteria/ Multidisciplinary Design*, 2004, pp. 1–40.

[12] Keane, A. J., and Nair, P. B., *Computational Approaches for Aerospace Design*, Wiley, New York, 2005.

[13] Ghisu, T., Parks, G. T., Jarrett, J. P., and Clarkson, P. J., "Adaptive Polynomial Chaos for Gas Turbine Compression Systems Performance Analysis," *AIAA Journal*, Vol. 48, No. 6, June 2010, pp. 1156–1170. doi:10.2514/1.J050012

[14] Carroll, T. C., and McAvoy, W. H., "The Formation of Ice upon Airplanes in Flight," NACA TN 313, 1929.

[15] Jacobs, E. N., "Airfoil Sections Characteristics as Affected by Protuberances," NACA TR 446, 1932.

[16] Noth, H., and Polte, N., "The Formation of Ice on Airplanes," NACA TM 786, 1936.

[17] Gulick, B. G., "Effects of Simulated Ice Formation on the Aerodynamic Characteristics of an Airfoil," NACA WR L-292, 1938.

[18] Gray, V. H., and Von Glahn, U. H., "Effect of Ice and Frost Formations on Drag of NACA 65-212 Airfoil for Various Modes of Thermal Ice Protection," NACA TN 2962, 1953.

[19] Bragg, M. B., Broeren, A. P., and Blumenthal, L. A., "Iced-Airfoil Aerodynamics," *Progress in Aerospace Sciences*, Vol. 41, No. 5, 2005, pp. 323–362. doi:10.1016/j.paerosci.2005.07.001

[20] Gray, V. H., "Correlations Among Ice Measurements, Impingement Rates, Icing Conditions, and Drag Coefficients for an Unswept NACA 65AO04 Airfoil," NACA TN 4151, 1958.

[21] Gray, V. H., "Prediction of Aerodynamic Penalties Caused by Ice Formations on Various Airfoils," NACA TN 2166, 1964.

[22] Potapczuk, M. G., and Gerhart, P. M., "Prediction of Aerodynamic Penalties Caused by Ice Formations on Various Airfoils," AIAA Paper 85-0410, 23rd Aerospace Sciences Meeting, Reno, NV, 1985.

[23] Kwan, O. J., and Sankar, L. N., "Numerical Study of the Effects of Icing on Finite Wing Aerodynamics," AIAA Paper 90-0757, 28th Aerospace Sciences Meeting, Reno, NV, 1990.

[24] Potapczuk, M. G., and Berkovitz, B. M., "An Experimental Investigation of Multi-Element Airfoil Ice Accretion and Resulting Performance Degradation," NACA TM 101441, 1989.

[25] Addy, H. E. J., Miller, D. R., and Ide, R. F., "A Study of Large Droplet Ice Accretion in the NASA Lewis IRT at Near-Freezing Conditions Part 2," NACA TM 107424, 1997.

[26] Lee, S., and Bragg, M. B., "Effects of Simulated-Ice Shapes on Airfoil: Experimental Investigation," AIAA Paper 99-0092, 37th Aerospace Sciences Meeting, Reno, NV, 1999.

[27] Chung, J., Reehorst, A., Choo, Y., Potapczuk, M., and Slater, M., "Navier–Stokes Analysis of Flowfield Characteristics of an Ice-Contaminated Aircraft Wing," *Journal of Aircraft*, Vol. 37, No. 6, 2000,

pp. 947–959.
doi:10.2514/2.2716
[28] Kumar, S., and Loth, E., "Aerodynamic Simulations of Airfoils with Upper-Surface Ice-Shapes," *Journal of Aircraft*, Vol. 38, No. 2, 2001, pp. 285–295.
doi:10.2514/2.2761
[29] Pan, J., and Loth, E., "Reynolds-Averaged Navier–Stokes Simulations of Airfoils and Wings with Ice Shapes," *Journal of Aircraft*, Vol. 41, No. 4, 2004, pp. 879–891.
doi:10.2514/1.587
[30] Pan, J., and Loth, E., "Detached Eddy Simulations for Iced Airfoils," *Journal of Aircraft*, Vol. 42, No. 6, 2005, pp. 1452–1461.
doi:10.2514/1.11860
[31] OpenFOAM User Guide, OpenCFD, http://www.openfoam.com/docs/user/index.php [retrieved 1 July 2009].
[32] Jasak, H., Weller, H., and Gosman, A., "High Resolution NVD Differencing Scheme for Arbitrarily Unstructured Meshes," *International Journal for Numerical Methods in Fluids*, Vol. 31, No. 2, 1999, pp. 431–449.
doi:10.1002/(SICI)1097-0363(19990930)31:2<431::AID-FLD884>3.0.CO;2-T
[33] Broeren, A. P., Lee, S., LaMarre, C. M., and Bragg, M. B., "Effect of Airfoil Geometry on Performance with Simulated Ice Accretion Volume 1: Experimental Investigation," NASA TR DOT/FAA/AR-03/64, 2003.
[34] Bragg, M. B., and Loth, E., "Effects of Large-Droplet Ice Accretion on Airfoil and Wing Aerodynamics and Control," Federal Aviation Administration TR DOT/FAA/AR-00/14, 2000.
[35] Lapworth, L., "Challenges and Methodologies in the Design of Axial Flow Fans for High-Bypass-Ratio, Gas Turbine Engines Using Steady and Unsteady CFD," *Advances of Computational Fluid Dynamics in Fluid Machinery Design*, Professional Engineering, Bury, St. Edmunds, U.K., 2003.
[36] Drela, M., "XFOIL 6.9 User Guide," TR, , Massachusetts Inst. of Technology, Cambridge, MA, 2001.
[37] Song, W., and Keane, A. J., "A Study of Shape Parameterisation Methods for Airfoil Optimisation," AIAA Paper 2004-4482, 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, New York, 30 Aug.–1 Sept. 2004.
[38] Samereh, J. A., "Survey of Shape Parameterization Techniques for High-Fidelity Multidisciplinary Shape Optimization," *AIAA Journal*, Vol. 39, No. 5, 2001, pp. 877–883.
doi:10.2514/2.1391
[39] Sederberg, T. W., and Parry, S., "Free-Form Deformation of Solid Geometric Models," *Computer Graphics*, Vol. 20, No. 4, 1986, pp. 151–160.
doi:10.1145/15886.15903
[40] Jaeggi, D. M., Parks, G. T., Kipouros, T., and Clarkson, P. J., "The Development of a Multi-Objective Tabu Search Algorithm for Continuous Optimisation Problems," *European Journal of Operational Research*, Vol. 185, No. 3, 2008, pp. 1192–1212.
doi:10.1016/j.ejor.2006.06.048
[41] Glover, F., and Laguna, M., *Tabu Search*, Kluwer Academic, Norwell, MA, 1997.
[42] Hooke, R., and Jeeves, T. A., "'Direct Search' Solution of Numerical and Statistical Problems," *Journal of the Association for Computing Machinery*, Vol. 8, No. 2, 1961, pp. 212–229.
doi:10.1145/321062.321069
[43] Kroo, I., "Aeronautical Applications of Evolutionary Design," Von Karman Institute Lecture Series on Optimization Methods and Tools for Multicriteria/Multidisciplinary Design, Rhode-Saint-Genèse, Belgium, 15–19 Nov. 2004.
[44] Kipouros, T., Jaeggi, D. M., Dawes, W. N., Parks, G. T., Savill, A. M., and Clarkson, P. J., "Biobjective Design Optimisation for Axial Compressors Using Tabu Search," *AIAA Journal*, Vol. 46, No. 3, 2008, pp. 701–711.
doi:10.2514/1.32794
[45] Ghisu, T., Parks, G. T., Jaeggi, D. M., Jarrett, J. P., and Clarkson, P. J., "The Benefits of Adaptive Parametrization in Multi-Objective Tabu Search Optimization," *Engineering Optimization*, Vol. 42, No. 10, Oct. 2010, pp. 959–981.
doi:10.1080/03052150903564882.
[46] Walters, R. W., and Huyse, L., "Uncertainty Analysis for Fluid Mechanics with Applications," NASA, CR 211449, 2002.
[47] Xiu, D., Lucor, D., Su, C.-H., and Karniadakis, G. E., "Spectral Representations of Uncertainty in Simulations: Algorithms and Applications," *Journal of Fluids Engineering*, Vol. 124, No. 1, 2002, pp. 51–59.
doi:10.1115/1.1436089
[48] Wiener, N., "The Homogeneous Chaos," *American Journal of Mathematics*, Vol. 60, No. 4, 1938, pp. 897–936.
doi:10.2307/2371268
[49] Lucor, D., Xiu, D., and Karniadakis, G., "Stochastic Modeling of Flow-Structure Interactions Using Generalized Polynomial Chaos," *Proceedings of the International Conference on Spectral and High Order Methods*, edited by C.-W. Shu, 2002.
[50] Xiu, D., and Karniadakis, G. E., "The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations," *SIAM Journal on Scientific and Statistical Computing*, Vol. 24, No. 2, 2002, pp. 619–644.
doi:10.1137/S1064827501387826
[51] Le Maître, O. P., Reagan, M. P., Najm, H. N., Ghanem, R. G., and Knio, O. M., "A Stochastic Projection Method for Fluid Flow: Part 1, Basic Formulation," *Journal of Computational Physics*, Vol. 173, No. 2, 2001, pp. 481–511.
doi:10.1006/jcph.2001.6889
[52] Monegato, G., *Fondamenti di Calcolo Numerico*, Edizioni Clut, Torino, Italy, 1998.
[53] Cameron, R. H., and Martin, W. T., "The Orthogonal Development of Nonlinear Functionals in Series of Fourier-Hermite Polynomials," *Annals of Mathematics and Artificial Intelligence*, Vol. 48, No. 2, 1947, pp. 385–392.
[54] Gautschi, W., "On Generating Orthogonal Polynomials," *SIAM Journal on Scientific and Statistical Computing*, Vol. 3, No. 3, 1982, pp. 289–317.
doi:10.1137/0903018
[55] Wan, X., and Karniadakis, G. E., "Multi-Element Generalized Polynomial Chaos for Arbitrary Probability Measures," *SIAM Journal on Scientific and Statistical Computing*, Vol. 28, No. 3, 2006, pp. 901–928.
doi:10.1137/050627630